

Moral Vagueness: A Dilemma for Non-Naturalism

Cristian Constantinescu

Pythagoreans, we are told, assigned the number 4 to justice. Probably parodying their claim, Socrates announces in the *Republic* that the difference between the just and the unjust man in respect of pleasure and pain is 729 (587b–588a). Whether intended or not, the effect of such claims is utterly comedic—almost on a par with declaring the meaning of life to be 42. By contrast, Aristotle sounds much more sensible when, reflecting upon the subject matter of his inquiry in the *Nicomachean Ethics*, he advises that “we must be content, in speaking of [fine and just actions, and of goods in general], to indicate the truth roughly, . . . for it is the mark of an educated man to look for precision in each class of things just so far as the nature of the subject admits” (1094b, 19–25). Common sense appears to side with Aristotle here: the very attempt to assign cardinal numbers to essentially imprecise, indeterminate values like justice, happiness, goodness, and the like, strikes most of us as deeply incongruous. Moral values belong to the throbbing centre of human life and, as Mark Sainsbury reminds us, “the throbbing centres of our lives appear to be describable only in vague terms” (1996: 251).

I shall take up here the particular type of indeterminacy invoked by Sainsbury: the vagueness of our moral terms. Predicates like “just,” “happy,” “cruel,” “generous,” “good,” “permissible,” etc. are vague in much the same way as “bald,” “tall,” “thin,” “red,” “heap,” and other paradigms of vagueness. Paradoxes aside, this phenomenon is, to my mind, mostly benign. Unlike other forms of indeterminacy, such as obscurity, ambiguity, or inaccuracy, vagueness does not greatly hinder moral inquiry. Rather, it can be taken to reflect a wholesome state of undecidedness concerning certain moral issues, indicating that we have left the door open for further debate and negotiation. More generally, acknowledging vagueness in a certain domain is often a sign that we recognize a diversity of human goals and interests, and

that we are prepared to tolerate variation of opinion and therefore a certain amount of disagreement within that domain.

The connection between vagueness and disagreement is easily brought out by reflecting on the phenomenology of borderline cases (one of the most salient symptoms of vagueness). In a borderline case, the semantic criteria governing a predicate's application seem to pull in opposite directions: if *a* is a borderline instance of the predicate "*P*", then there are strong reasons in favour of classifying *a* as *P*, and strong reasons in favour of classifying *a* as not-*P*. The clash between such reasons can be experienced internally by one speaker, leading to perplexity or ambivalence,¹ or externally by two or more speakers, leading to disagreement.

This interesting link between vagueness and disagreement has not gone unnoticed in the literature. Aristotle himself, in the passage quoted earlier, explains the relative imprecision of ethics by appealing to the fact that "noble and just actions [and all goods, more generally] exhibit much variety and fluctuation" (*Nicomachean Ethics* 1094b, 15–16). And closer to our time, a number of philosophers seeking to defend moral realism against the so-called "argument from disagreement" (namely, the charge that realists cannot account for the persistence and pervasiveness of seemingly intractable moral dissensions) have pointed out that ethical objectivity is compatible with the existence of pockets of indeterminacy in our moral frameworks, which might in turn explain the occurrence of disagreement.² The general strategy of this defensive move (which I shall call the "vagueness defence") is well sketched in the following quote from Russ Shafer-Landau, who has also provided the most sustained examination of moral vagueness in the literature so far:³

I suggest that the apparent soundness of the central noncognitivist argument [from disagreement] may be dispelled if we abandon an assumption long associated with objectivism, namely, that morality is entirely determinate. . . . If we allow for moral indeterminacy . . . , then we have a promising explanation of disagreement for the objectivist. In those situations where perfect unanimity seems a pipe dream, the objectivist can attempt to show that an ineliminable element of moral indeterminacy exists for the situation being debated. The fact that there is no uniquely correct assessment awaiting discovery can appropriately explain why in some cases even idealised agents would fail to converge on the identity of a single best moral evaluation. (1994: 332, 336)

¹ Ambivalence, as I use it here, is thus a form of inner disagreement, where the speaker is pulled in opposite directions. We often express this by saying things like "I'm conflicted," "I'm torn," or "I'm in two minds about it."

² See Hurley 1992; Parfit 2011: ii. 559–62; Railton 1992; Shafer-Landau 1994; Sosa 2001; Vasile 2010; Wolf 1992.

³ See in particular Shafer-Landau 1995.

Recently, Derek Parfit has also concurred:

Some questions may be *indeterminate*, in the sense that they have no answer. That is sometimes true, for example, of the question “Is he bald?” If some man has no hair, he is bald. If some man has a full head of hair, he is *not* bald. But we cannot plausibly assume that, in all cases between these two extremes, any man must either *be*, or *not* be, bald. In many cases, though it is not true that some man is bald, it is also not true that this man is *not* bald. Similar claims might apply to normative questions. . . . If some normative questions are indeterminate, having no answer, this would provide another explanation of some normative disagreements. When people disagree about whether some act is wrong, they may mistakenly assume that this act must either be, or not be, wrong. If these people gave up this assumption, they might often cease to disagree. (Parfit 2011: ii. 559–60, 562)

On the face of it, this looks like a plausible suggestion. There are, as we have seen, strong conceptual links between vagueness and disagreement, so the vagueness defence promises to defuse the argument from disagreement: the realist can acknowledge the possibility of intractable moral disputes, but insist that they often arise from the vagueness of our moral terms.

Yet, despite the initial plausibility of the vagueness defence, I shall argue in what follows that not all moral realists can readily employ this move. On close examination, vagueness turns out to be incompatible with tenets that are key to certain forms of moral realism. My primary target here will be the non-naturalist realism defended by philosophers like Shafer-Landau and Parfit. I aim to show that by invoking moral vagueness in response to the argument from disagreement, non-naturalists get embroiled in a dilemma, either horn of which forces them to give up some of their central commitments.

1. A SKETCH OF THE MAIN ARGUMENT

For simplicity, I shall use the label “non-naturalism” for the brand of realism championed by Shafer-Landau.⁴ I take the following to be key tenets of this view, which will be relevant to my argument:

⁴ I focus on Shafer-Landau because he most explicitly upholds both non-naturalism and the vagueness defence. Parfit, as we have seen, is also a good example of a non-naturalist who appeals to moral vagueness, but the extent to which he would assent to all of the seven theses I discuss here remains slightly unclear (see n. 5). More widely, non-naturalism seems to have enjoyed a vigorous revival recently, as many of the current leading metaethicists have embraced its main tenets in one form or another: see e.g. Audi 2004; Crisp 2006; Cuneo 2007; Dancy 2006; Enoch 2011; FitzPatrick 2008; Huemer 2005; Wedgwood 2007.

Cognitivism: Moral sentences express beliefs and are therefore truth-apt.

Correspondence: Moral sentences are true when they correspond to moral facts.

Atomism: Moral facts are instantiations of moral properties.

Objectivism: Moral facts and properties are mind-independent.

Supervenience: Moral facts and properties supervene upon natural facts and properties.

Non-reductivism: Moral facts and properties are metaphysically *sui generis*.

Rationalism: Moral facts are intrinsically reason-giving.

All these are claims that Shafer-Landau endorses explicitly, and defends forcefully (see Shafer-Landau 2003). I shall have more to say about some of them in the course of my argument. For now, let me just point out that the view that emerges by putting together such claims is a *robust* form of realism: the non-naturalist I have in mind will not adopt a minimalist conception of moral properties and facts, and will not content himself with characterizing them merely as “those things, whatever they are, which are picked out by our true moral judgements”; nor will he embrace a quietist stance, downplaying the ontological commitments of talk of moral facts and properties (à la Scanlon 2003 or Dworkin 2011).⁵ Rather, the earnest non-naturalist will insist that moral properties are qualities which

⁵ Parfit sometimes seems to adopt this strategy too in his (2011). Though he seems quite happy to countenance talk of non-natural, irreducibly normative properties at various points throughout the book, in §112 he puts forward a view he terms “non-metaphysical cognitivism,” according to which: “There are some claims that are irreducibly normative in the reason-involving sense, and are in the strongest sense true. But these truths have no ontological implications. For such claims to be true, these reason-involving properties need not exist either as natural properties in the spatio-temporal world, or in some non-spatio-temporal part of reality” (2011: ii. 486). Parfit’s “non-metaphysical” view relies on distinguishing between an ontological reading of “exist” (on which the claim that normative properties exist comes out false), and a non-ontological reading (on which the claim is true). I’m not sure a view that appeals to a notion of “non-ontological existence” is properly characterized as “non-metaphysical,” so much as “super-metaphysical.” Perhaps the best way to understand Parfit’s suggestion is as a kind of quietism about metaphysical matters. But then it becomes difficult to understand how any substantive debate between naturalism and non-naturalism, of the kind Parfit himself happily engages in, can even be had. If, as Parfit appears to suggest at times, non-naturalism is to be characterized simply as the claim that there are irreducibly normative truths and concepts, then his view becomes indistinguishable from the non-analytical naturalism defended by philosophers like Gibbard (2006), according to which, although there are irreducibly normative claims, there are no irreducibly normative facts, because normative concepts and natural concepts signify properties of the same kind. Since Parfit opposes Gibbard’s non-analytical naturalism, he must think that not just moral claims, but moral properties too, are irreducibly normative. To then pull a Meinongian stunt and add that

individuals can instantiate (as opposed to being just sets of individuals, for instance), and which, while “resulting from,”⁶ “being realised by,” or “being constituted by”⁷ natural properties, are nevertheless neither identical nor reducible to natural properties.

I am going to present my argument against this type of earnest non-naturalism in the form of a dilemma. According to the non-naturalist, there are *sui generis*, irreducible moral properties. But if such properties exist, then given the vagueness of many of our moral predicates, either

Horn 1: vague moral predicates pick out vague moral properties,

or

Horn 2: vague moral predicates pick out sharp moral properties, and vagueness arises from a different source.

I will argue that by taking Horn 1 of this dilemma, the non-naturalist ends up with moral properties that are either mind-dependent or reducible, in ways that are incompatible with either Objectivism, Supervenience, or Non-reductivism above. On the other hand, by taking Horn 2 the non-naturalist can avoid mind-dependence and reducibility, but only by countenancing instead an ontology of perfectly sharp, strongly unknowable moral properties, which I shall claim ultimately undercuts his commitment to either Rationalism or Supervenience. Either way, I will conclude, the non-naturalist cannot accommodate moral vagueness without giving up some of his key commitments.

2. FIVE REASONS TO BELIEVE IN MORAL VAGUENESS

So far, I have simply assumed without argument that moral predicates are vague. While I find this claim absolutely intuitive, some philosophers don't.⁸ Although my argument targets non-naturalists who, like Shafer-Landau,

such properties exist only in a non-ontological sense appears to me to nullify the whole debate. The moral of the story is that the earnest non-naturalist must take moral properties seriously, as of course Shafer-Landau and the other non-naturalists listed in n. 4 all do.

⁶ See Dancy 1981. Dancy takes his notion of “resultance” from Ross 1930, but significantly sharpens it. I should note that Dancy's view at this early stage wasn't decidedly non-naturalist, as he allowed the compatibility of resultance with a relation of token-identity between normative and natural properties.

⁷ See Shafer-Landau 2003: 72–9.

⁸ Dworkin e.g. has long argued that every legal question must have a perfectly determinate answer (Dworkin 1977), and has recently extended this view to the ethical domain (Dworkin 2011: ch. 5).

believe in moral vagueness, I also hope to convince those who, like Dworkin, find this idea doubtful. My aim in doing this, of course, is to weaken the non-naturalist position more widely: my dilemma, I submit, applies not just to non-naturalists who are already committed to moral vagueness, but to all non-naturalists, since we are all committed to moral vagueness. I shall therefore give five interrelated reasons for thinking that at least some (and probably most) of our moral predicates are vague. Due to their interrelatedness, these reasons can be expressed as a single claim: moral predicates display all the standard symptoms of vagueness: (a) imprecise gradability; (b) boundarilessness; (c) borderliness; (d) tolerance; (e) soriticality.

- (a) *Imprecise gradability.* Like vague predicates from other domains of discourse, moral predicates are imprecisely gradable. Their gradability means that they admit of degrees: one person can be *very* honest, another *slightly* cruel, a third *not too* courageous. Furthermore, one person can be *more just* than another, and often *by much*. But, *pace* Socrates, it seems impossible to tell precisely by *how* much. The reason for this is that justice, honesty, cruelty, courage, and all the other moral values, simply do not appear to be the kinds of properties that could come in discrete, cardinally quantifiable units. I must stress that the focus here is on *imprecise gradability*, rather than simply impreciseness or gradability taken separately. Famously, impreciseness is insufficient for vagueness: “natural number between 1 and 100” is imprecise, but not vague. And where a property is precisely gradable, vagueness can, but does not necessarily, arise: both “tall” and “between 100 and 200 cm. high” pick out properties which are precisely gradable (heights), but the former is vague while the latter is not. This raises an interesting point about the properties denoted by vague predicates: in many cases, these properties supervene upon properties which are precisely gradable, but are not themselves precisely gradable. Consider the predicate “bald.” Baldness supervenes on the property of having x hairs on one’s scalp, which is precisely gradable. But baldness itself is not precisely gradable: even if we know precisely how many more hairs Abe has on his scalp than Ben, it doesn’t follow that we thereby can tell how much balder Ben is (if for no other reason than at least because other factors, like the distribution of hair on one’s scalp, contribute to our ascriptions of baldness, thereby making gradability imprecise). Similarly for moral properties: honesty supervenes on the number of lies one tells. But it doesn’t follow that if we know how many lies Abe and Ben have told so far in their lives, we can thereby know how much more honest the one is than the other.
- (b) *Boundarilessness.* Imprecise gradability leads to imprecise boundaries. Because honesty does not come in precise degrees, there cannot be a

fixed degree of honesty constituting the threshold for whether someone counts as honest or not. Thus, the predicate “honest” fails to separate sharply between individuals to which it applies (its extension) and individuals to which it doesn’t (its anti-extension). Suppose, for simplicity, that the number of lies told were the only criterion for classifying people as honest (while gravity, context, intentions, and countless other factors, didn’t matter). Imagine arranging people in a continuum, based on how many lies they told: it seems implausible to suggest that there will be a precise cut-off point (one white lie) separating the honest from the dishonest.

- (c) *Borderlineness*. From boundarilessness we get to borderlineness. When a predicate lacks a precise boundary between its extension and its anti-extension, it also admits of borderline instances: individuals to which the predicate neither determinately applies, nor determinately doesn’t apply. Along the continuum we have just imagined, there will be people in the fuzzy area between the extension and the anti-extension of “honest”. These people are borderline instances of the predicate: not clearly honest, but not clearly dishonest either.
- (d) *Tolerance*. Like paradigmatically vague predicates, many moral predicates are “tolerant” with respect to small changes in the properties on which their application is based. Consider Abe, a determinately honest man: Abe may have told a few lies in his life, but they were mostly intended to save others from embarrassment, or forms of being polite. When it comes to the important things, Abe is always truthful. Now add one very small lie to Abe’s history. This seems insufficient to turn Abe from a clearly honest into a clearly dishonest person. In this sense, the predicate “honest” is tolerant with respect to small lies.
- (e) *Soriticality*. Finally, tolerance famously gives rise to sorites arguments: if Abe counts as honest having told just three small lies in his fairly long lifetime so far, then so does Abe after telling four small lies; but if Abe counts as honest after telling four small lies, then so does Abe after telling five lies, and so on until we get the absurd conclusion that Abe is honest despite having told a million lies.

Readers with strong Kantian inclinations may find this example unconvincing. They might think, for instance, that there is a precise cut-off point between being honest and not being honest: namely, the very first lie. The duty not to lie is a perfect duty in Kant’s system, and that gives us an absolute threshold. Nevertheless, examples of a similar kind can be generated with ease. If causing slight pain (say, a barely perceptible pinprick) when giving someone an injection isn’t cruel, then causing ever-so-slightly more

intense pain isn't cruel either; but applying tolerance a sufficient number of times would yield the absurd conclusion that causing agonizing pain isn't cruel. If watching hard porn isn't harmful to Ben today, as he turns 40, then watching hard porn wasn't harmful yesterday; apply tolerance over and over again, and eventually you get the conclusion that watching hard porn wouldn't have been harmful to Ben when he was 7.⁹ And lest one should think that only thick moral concepts can be vague, consider the case of "wrong" in the following scenario (borrowed from Sorensen 1990): it's definitely wrong to draw 5,000 millilitres of blood from a person for blood tests; if so, then it's also wrong to draw 4,999 millilitres of blood; apply this enough times, and you get the absurd conclusion that it's wrong to draw 1 millilitre of blood. Finally, think of "ought" in this scenario: I ought to give £10 to charity every month; if so, then I also ought to give £10.01; but if I apply tolerance enough times, I get the absurd conclusion that I ought to give all my money to charity.¹⁰

3. THE DILEMMA UNPACKED

I hope this is enough to convince readers that moral predicates can be vague. I turn now to the task of expounding my dilemma. If the argument is to stick, I must show that each horn is problematic for the non-naturalist. It will help if we fix our attention on an example. It's Saturday morning and Abe, who is a psychotherapist, is at home relaxing with his family. It's been a while since he has had the chance to do so, for work has been quite hectic lately. Abe takes genuine pleasure in being at home with his family, and he's also promised he wouldn't sacrifice family time for work this weekend. However, Abe's phone rings: it's his patient, Ben, who sounds very distraught and desperately needs to talk to him, or else "he'll do something stupid." Abe knows that Ben has been very unstable recently, and so decides to go and meet him at his office. It seems clear, given Ben's state and despite Abe's promise to his family, that in this case it's permissible for Abe to go talk with Ben for an hour. On the other hand, spending the entire day with Ben would clearly not be permissible: not only has Abe made a promise to his family, but there are other therapists at hand who could take over,

⁹ Of course, in a case like this, the law does set an arbitrary cut-off point at the age of 18. But no one seriously thinks that this removes the vagueness of "harmful": if watching hard porn one second before midnight on the day of his 18th birthday is harmful to Ben, so is watching hard porn one second later (despite the latter act being legally permitted, and the former not).

¹⁰ I thank an anonymous reviewer for pressing me on the need to refine my examples of moral vagueness.

not to mention that Ben should not be encouraged to develop an emotional dependency on Abe. Let us say, then, that it's clearly impermissible for Abe to spend eight hours with Ben. Now, somewhere between one hour and eight hours along this continuum, there will be a fuzzy area containing actions (such as Abe's spending 120 minutes with Ben, Abe's spending 121 minutes, and so on) which are neither clearly permissible, nor clearly impermissible.

Eventually, Abe spends two hours with Ben (call this action of his "*A*"). Let us suppose that in this case *A* is a borderline instance of the predicate "permissible." According to the non-naturalist, the predicate "permissible" denotes an objective, irreducibly normative property, PERMISSIBILITY.¹¹ What should the non-naturalist say about the property of PERMISSIBILITY itself, when confronted with a borderline instance of it? It seems to me there are only two options:

- (i) *A* is a borderline instance of "permissible" because PERMISSIBILITY is vague;
- (ii) *A* is a borderline instance of "permissible" despite PERMISSIBILITY being sharp.

This generalizes to all moral properties. *Vis-à-vis* any moral statement containing a vague moral predicate, the non-naturalist can either choose to treat that predicate as denoting a vague moral property, or else insist that all moral properties are sharp, and consequently seek to explain predicate vagueness in some other way. Thus we get the two horns of my dilemma: the non-naturalist can opt for Horn 1 by committing himself to an ontology containing vague moral properties, or for Horn 2 by countenancing only sharp moral properties. In virtue of his allegiance to robust moral properties, the non-naturalist cannot avoid this choice. But I will argue that each horn leads to serious problems. Such problems, I believe, can only be avoided by banishing moral properties altogether.

Before I go on to examine each horn in detail, a clarification is required. It may seem that by asking the non-naturalist to opt for either a vague or a sharp moral ontology, I am in fact arbitrarily forcing a choice between an *ontic* conception that would locate vagueness "out there," in the moral properties, and an *epistemic* conception, which would insist that moral properties are perfectly sharp and regard vagueness instead as the result of imperfect knowledge. It may thus seem odd that the problem is framed as a dilemma, rather than as a *trilemma*. After all, why would the non-naturalist not be entitled to a *semantic* account of moral vagueness? This challenge,

¹¹ From here on I adopt the convention of using small capitals to signify properties.

however, misconstrues the nature of my dilemma. The choice I am trying to force is based on the ontological issue of whether moral properties should be taken as vague or sharp. On that issue, supporters of onticism and epistemicism are indeed sharply divided, each being forced to take either one or the other horn. The semanticist, however, can go either way. Like Russell (1923) or Lewis (1986a), she may insist that the vagueness infecting our language is entirely due to the general indeterminacy of our representations: the world itself is fully determinate. In terms of our dilemma, this would amount to choosing Horn 2. Alternatively, like Shapiro (2006) and Hyde (2008), the semanticist may instead deplore Russell's insistence that the world must be crisp as a mere reflex of what Whitehead called "the myth of independent existence": the idea that the world can be conceived of in separation from our representations. Instead, this latter kind of semanticist would regard ontic vagueness as a necessary accompaniment to semantic vagueness.

The picture we should expect is then this: on Horn 1 of the dilemma we'll find both advocates of onticism and some supporters of semanticism, countenancing vague moral properties but offering competing accounts thereof; on Horn 2, other semanticists will agree with epistemicists that moral properties are sharp, but disagree about the explanation of predicate vagueness. Analysing each horn carefully will require paying close attention to these subtle differences.

3.1. Horn 1: Vague Moral Properties

The non-naturalist gets into Horn 1 of my dilemma by choosing to account for moral vagueness in terms of vague moral properties.¹² Various accounts of vague properties have been proposed in the general literature on vagueness (Sainsbury 1989; Tye 1990; Rosen and Smith 2004; Shapiro 2006; Hyde 2008; Schiffer 2010). As a first approximation, we may start by characterising a property as vague just in case it admits of borderline instances. Thus, BALDNESS is vague because there are people who are borderline bald, and in our example above PERMISSIBILITY is vague because it is indeterminate whether Abe's action *A* (spending 2 hours with Ben) instantiates it. Some authors are in fact content with this minimal account of vague properties (Sainsbury 1989; Tye 1990; Hyde 2008). But the initial characterization doesn't suffice. For one thing, a property may have a borderline instance despite being perfectly sharp: for instance, if its borderline

¹² This is indeed Shafer-Landau's strategy: in his (1994) he argues for the "worldly" indeterminacy of moral properties, and against semantic and epistemic accounts.

instance is a vague object.¹³ Moreover, since *A*'s being a borderline instance of PERMISSIBILITY means that it is indeterminate whether *A* instantiates PERMISSIBILITY, the attempt to define vague properties in terms of borderline instances seems guilty of a fallacy: it moves from "It is indeterminate whether *A* instantiates PERMISSIBILITY" to "Of the property PERMISSIBILITY, it is indeterminate whether it is instantiated by *A*."¹⁴ And finally, if the minimal characterization was all we had, we couldn't even make sense of the differences between ontic, semantic, and epistemic theories of vagueness. For arguably, all of these theories can accept the truism that some properties have borderline instances, but not all of them would countenance vague properties. Thus, if it makes sense to enquire whether there really are vague properties, then the debate between supporters of onticism, semanticism, and epistemicism must be a debate over vague properties in a more robust sense.

As I have indicated, both onticists and certain kinds of semanticists can accommodate talk of vague properties. Let's take these views in turn, starting with semanticism.

3.1.1. *The Semantic View of Vague Properties*

The first theory that springs to mind under the heading of "semanticism" is surely supervaluationism. Yet, as we shall shortly see, supervaluationists take the properties picked out by vague predicates to be perfectly sharp, and for this reason we shall discuss their views when analysing Horn 2 of the dilemma. For now, we must look elsewhere for a semantic account of vague properties.

Besides supervaluationism, the other major brand of semanticism about vagueness is contextualism. I shall use Shapiro's (2006) contextualist account of vagueness for illustration, as it explicitly countenances vague properties. Shapiro takes properties in general to be individuated by the word-usage practices of competent speakers with regard to the predicates that express them. A useful model for understanding these linguistic practices is to think of them as conventions established by conversational partners who keep logs of their conversations (this is inspired by Lewis's notion of a "conversational score"). There are certain words whose application the conversationalists always agree on. For instance, in the case of the predicate "natural

¹³ Thus, it may be indeterminate whether a particular cloud weighs 1 billion kilograms not because of any vagueness in the property WEIGHS 1,000,000,000 KG, but because the cloud itself lacks sharp boundaries.

¹⁴ As Williamson points out in a similar connection, this is no more valid than the fallacious move from "It is contingent whether the number of planets is even" to "Of the number of planets, it is contingent whether it is even" (see Williamson 2003: 701–2).

number,” competent speakers will be able to decide, of any object they are presented with, whether or not that object is a natural number. A determinate predicate individuates a determinate property: thus, the property of being a natural number is sharp. On the other hand, there are also cases where conversational scores may remain silent: the conversationalists’ previous word-use practices have not established conventions that completely determine any future applications of such predicates. The predicate “bald” is a typical example. We can safely assume that the linguistic conventions governing the application of this predicate have not settled criteria of correct application for every conceivable configuration of hair on a human being’s scalp. Therefore, there will be unclear cases, over which competent speakers might disagree. There will also be cases about which they will feel ambivalent: one and the same speaker may classify me as “bald” in a context where the standard of hirsuteness is set in relation to Jimi Hendrix, and as “not bald” in a context where I’m standing next to Billy Corgan. All of this is, of course, very familiar. The point is that, on this account, properties are individuated by the semantic criteria governing the use of predicates denoting them, which are in turn fixed by conventions made by speakers.¹⁵

This rough characterization should suffice to suggest that the contextualist approach to vague properties is not an appealing option for the non-naturalist who wants to account for moral vagueness by appealing to vague moral properties. On Shapiro’s account, vague properties turn out to be *judgment-dependent* in a deep sense: not only are they *responsive* to certain linguistic/psychological facts about the competent speakers—they are actually *constituted* by such facts. The analogue of this in the moral realm would be a form of constructivism, grounding moral properties in the practices of rational agents negotiating rules for mutual behaviour. But on this view, speakers confronted with borderline instances of moral predicates cannot be assumed to *track* an independent realm of properties in their judgements: the direction of fit is Euthyphronic (from speakers to properties) rather than Socratic (from properties to speakers). Adopting such an account of vague properties would therefore seem to be incompatible with the thesis of Objectivism.

But perhaps we’ve gone too fast. There is an obvious rejoinder to the preceding line of argument: while it may be true that vague properties are in

¹⁵ This approach to vague properties may be supplemented with a psychological account (in the manner of Schiffer 2000). After all, as Shapiro himself acknowledges (2006: 24ff.), the picture of communicators keeping conversational scores appears to leave one important question unanswered: what makes it the case that individual ambivalence or collective disagreement are sometimes appropriate attitudes for the conversationalists to display? Schiffer’s account may offer the key, by explaining vagueness in terms of belief-degrees.

one sense judgement-dependent, this may only be the case *within* their borderline areas. Determinate instances, on the other hand, could be regarded as mind-independent. This is actually Shapiro's view:

With vague predicates, judgment-dependent matters sometimes figure in how the extension is fixed. For determinate cases, the judgment-dependent matters in question are trumped by other factors that fix the extension. In the borderline region, the judgment-independent features give out, and, consequently, the judgment-dependent features of the process dominate. The predicate becomes Euthyphronic. As I see things, every vague predicate has (or can have) such a Euthyphronic region. But this does not undermine the overall objectivity of the predicate, or the discourse in which it figures. If it did, there would be precious little objectivity anywhere. (2006: 209)

If Shapiro's remarks here are correct, perhaps the non-naturalist can embrace the preceding account of vague properties after all.

To see whether this will work, let us first translate Shapiro's talk of vague predicates in this passage into talk of vague properties. Take "bald," which Shapiro uses as an example of a predicate with both a Socratic and a Euthyphronic region. Presumably this means that the property BALDNESS denoted by this predicate also has such "regions": a mind-independent core, determined entirely by objective factors, and a mind-dependent penumbra, determined at least in part by speakers' conventions. The first image that comes to mind is that of a fuzzy set of bald individuals, with different degrees of membership: some individuals are determinately bald and therefore full members, while others are indeterminately bald and therefore members to lesser degrees. For the former kind of individuals, objective factors fix their membership status, whereas for the latter it is up to speakers to decide in each case whether it is appropriate or not to count them as members of the set. Clearly, this isn't an understanding of vague properties that the non-naturalist can readily embrace and apply to moral properties. For the non-naturalist, the vague property of PERMISSIBILITY, for instance, is more than just a set of actions, for there is nothing metaphysically *sui generis* or intrinsically reason-giving about sets.

An alternative is to think of the different "regions" of a vague property more literally, i.e. mereologically. Thus, a vague property could be more like a cloud than a set. A cloud is made up of a core of particles definitely within it, and a looser, marginal region of particles of which it's indeterminate whether they are parts of the cloud. At first blush, this model seems more compatible with Shafer-Landau's conception of moral properties as being "constituted by," though not identical to, natural properties (Shafer-Landau 2003: 72–9). PERMISSIBILITY, for instance, could be a non-natural type-property made up of various token-properties (PERMISSIBILITY-instantiations, or tropes), each

constituted or realized entirely by clusters of natural tropes.¹⁶ The difference between tropes belonging to the property's core and tropes in the property's penumbra could then be understood in terms of a further type of natural fact belonging to the constitution of the latter, but not the former: namely, natural facts about speakers' linguistic conventions/attitudes. Due to being constituted in part by such natural facts about speakers, tropes in the penumbra of PERMISSIBILITY could be considered mind-dependent in the Euthyphronic sense. However, the tropes in the property's determinate core would remain perfectly objective.

But there is a complication for this view. Ascriptions of vague predicates are notoriously context-relative: as noted earlier, I can count as determinately not-bald when compared to Billy Corgan, but indeterminately bald when compared to Jimi Hendrix. The same applies, of course, to vague moral predicates: I may count as indeterminately generous in many contexts, but I'm definitely not generous when compared to Peter Singer. Similarly, one and the same action can be definitely permissible when compared to one set of alternatives, and indeterminately permissible when compared to another set. On the picture we are now contemplating, this means that one and the same permissibility-trope, call it " P_i ", is part of the property's core from one perspective, and part of the property's penumbra from another. This in turn implies that P_i is constituted in one context by a set of natural tropes, call it " N_i ", and in another context by $N_i \cup \{s1, \dots, sn\}$, where " $s1$ ", \dots , " sn " denote natural facts about speakers' conventions and/or attitudes. However, this violates a fundamental constraint on constitution. As Ridge (2007: 342) points out, if constitution is to do the job Shafer-Landau wants it to—namely, that of accounting for Supervenience—then two conditions must obtain: (i) if a cluster of natural tropes N_i necessarily constitutes a moral trope P_i , then whenever N_i is present, P_i is present too; and (ii) if N_i constitutes P_i , then N_i necessarily constitutes P_i . But it's quite clear that these two conditions cannot be met for tropes of vague properties: due to the contextual variation noted above, one and the same moral trope will in some contexts belong to the core, and in other contexts to the penumbra of a moral property. But then it follows, on the present understanding of penumbral instances, that one and the same moral trope can be instantiated by different natural tropes in different contexts (since facts about speakers will sometimes feature, and sometimes not, in the trope's constitution). And this violates (i) above, thereby putting the claim of Supervenience in jeopardy.

¹⁶ This is in line with Ridge's reconstruction of Shafer-Landau's conception of constitution (Ridge 2007: 340ff.)

No such difficulties would arise, of course, for a moral naturalist. The naturalist could embrace Shapiro's conception of vague properties, while identifying vague moral properties with fuzzy sets of individuals, in the manner contemplated above. Non-naturalism, on the other hand, appears incompatible with a semantic view of vague properties.

3.1.2. *The Ontic View of Vague Properties*

The alternative approach to vague properties aims to give a purely *structural* account in terms of the formal characteristics distinguishing vague from sharp properties,¹⁷ rather than formulate individuating conditions in terms of *constitutive* facts about speakers' linguistic conventions or mental states.

This starts from the observation, already adumbrated (see n. 13), that there are two ways in which a property can have borderline instances: either when the property itself is vague, or when its instances are vague objects. Consider now a sharp property: if such a property has borderline instances, that can only happen because the objects in those instances are themselves vague. To take our earlier example, since WEIGHS 1,000,000,000 KG is a sharp property, only a vague object (e.g. a cloud) could be a borderline instance of it. In other words, sharp properties are "the properties that make vague objects out of their vague instances" (Rosen and Smith 2004: 187). This insight offers us a key: we could start by zeroing in on sharp properties, and then proceed by defining vague properties negatively (as those properties which are not sharp).

To illustrate further, take the vague property TALLNESS, and suppose Bill is a borderline instance of it. Must there be any vagueness about Bill himself? Not necessarily.¹⁸ It is possible for Bill to be a perfectly determinate body, yet a borderline instance of "tall." Compare that to the predicate "precisely 186 cm tall." Measurement imprecision aside, it is impossible for Bill to be a borderline instance of this precise predicate without being vague himself (perhaps, for instance, there's a hair loosely attached to his head in such a way that it is indeterminate whether or not it is part of his body, and counting that hair may alter our measurement). Rosen and Smith's own example involves a comparison between the vague colour-predicate "blue" and the precise predicate "blue-17", which denotes a particular point-sized region in the Munsell Colour Solid. Again, an object can be a borderline instance of the latter predicate only by being vague itself.

¹⁷ See Rosen and Smith 2004. The following paragraphs are based (at times only loosely) on their approach, adding qualifications drawn from Sanford, who offered similar suggestions much earlier (Sanford 1966, 2013).

¹⁸ Though of course there may well be, in the sense of there being at least one particle x such that x is neither determinately inside Bill, nor determinately not inside Bill.

Rosen and Smith distinguish between vague and sharp properties by generalizing this idea. On their account, vagueness turns essentially on issues of specificity. Properties are sharp if they are *maximally specific*—much like determinate point-sized regions in the colour solid, or point-sized units along a continuum of heights. For this reason, Rosen and Smith also refer to sharp properties as “point-properties.” Vague properties, on the other hand, lack specificity. Rather, they cover ranges of specific properties—much like “blue,” which covers a range of different shades, or “tall,” which covers a range of different heights. For this reason, we may call them “range-properties.”¹⁹

The distinction looks neat, but it does require an important qualification. Rosen and Smith appear to overlook the fact that there are properties which are sharp, but not maximally specific (i.e. not point-properties): e.g. the properties denoted by “taller than 186 cm,” or “taller than 186 cm but not taller than 189 cm.” These are imprecise, but sharp (i.e. not vague). That some range-properties turn out to be sharp may seem to throw some doubt on the usefulness of Rosen and Smith’s distinction. But the issue can easily be fixed, by noting a common characteristic of such sharp range-properties: they cover *precisely bound* ranges. This enables us to revise the above characterization of vague properties, so that it now reads: vague properties are properties that cover *indeterminate* or *unbound* ranges of point properties.

There are good reasons to think that moral properties are vague in the sense just specified. Like colour terms, moral predicates are general enough to apply not just to singular instances, but to entire ranges of cases. Moreover, these ranges are never precisely bound: there are no minimal or maximal values of COURAGE, KINDNESS, PERMISSIBILITY, etc., as shown in §2 above. And importantly, from a metaphysical point of view this structural account of vague properties seems less threatening than the semantic view examined in §3.1.1, since it doesn’t rely on facts about speakers’ conventions which could introduce an element of subjectivity. It might seem as if this ontic view of vague properties could be much more promising for the non-naturalist seeking an account of moral vagueness.

Under closer scrutiny, however, things appear more problematic. On the present account, the lack of specificity displayed by vague properties means that they are somehow *composite*, or *structured*: as we have seen, they cover fuzzy ranges, and as such are comprised of maximally specific

¹⁹ This could also be paraphrased in terms of a distinction between determinables and determinates: range-properties are determinates and determinables, whereas point-properties are perfect determinates (i.e. determinates but not determinables). For more on this, see Sanford 2013.

point-properties. This encourages the view that they are somehow *conceptually derivative*: i.e. capable of being broken down and understood in terms of sharp ones. For notice that corresponding to any vague property P there will be a set of properties $\{Q_1, Q_2, \dots\}$ such that each Q_i belonging to that set will be (i) perfectly sharp, and (ii) necessarily sufficient for P . But then it follows that the disjunction of all these sharp properties, $Q_1 \vee Q_2 \vee \dots$, will be necessarily coextensive with P . This seems intuitive: when I say that a is blue, I locate a within a fuzzy region on the colour spectrum, while leaving it open which particular shade of blue a may be; thus, my claim can be equated to “ a is $B-1$ or $B-2$ or $\dots B-17$ or \dots ”²⁰

The upshot is that vague properties are essentially *disjunctive*: more specifically, they can always be resolved into disjunctions of sharp properties. This is apt to spell trouble for the moral naturalist contemplating an ontology of irreducible moral properties. For one thing, the very notion of a “disjunctive property” has seemed incoherent to many philosophers (e.g. Lewis 1986b; Armstrong 1978). Here’s how Armstrong puts the issue:

Disjunctive properties offend against the principle that a genuine property is identical in its different particulars. Suppose a has a property P but lacks Q , while b has Q but lacks P . It seems laughable to conclude from these premisses that a and b are identical in some respect. Yet both have the “property”, P or Q . (1978: 20)

But perhaps not all disjunctive properties are as gerrymandered as Armstrong supposes they must be. There is a difference, for instance, between the following (putative) properties: PALE YELLOW OR BRIGHT YELLOW; PALE YELLOW OR DARK GREEN; YELLOW OR ANGRY (Sanford 1970, 2013). While the last of these definitely falls prey to Armstrong’s criticism, it may seem doubtful whether the second does, and quite certain that the first doesn’t. The issue, it may be thought, turns on *resemblance*. At least the first of these properties can be regarded as a real, genuinely disjunctive property, because there are deep similarities between its disjuncts: if a is pale yellow and b is bright yellow, then a and b are indeed “identical in some respect,” and therefore can be properly said to satisfy the predicate “pale yellow or bright yellow.”

This may seem to provide a satisfactory response to Armstrong’s attack. After all, most of the vague predicates in our language (“red,” “tall,” “bald,” but also “generous,” “honest,” “just,” “courageous,” etc.) appear to “carve nature at its joints” rather than just being artificially gerrymandered disjunctions (in other words, they are more like “pale yellow or bright yellow” than like “yellow or angry”). “Red,” for instance, is equivalent to

²⁰ Not in the sense that I must mean or intend that disjunction when I ascribe blueness to a , of course. The claim concerns the extensions of vague properties, not their intensions.

an indeterminate disjunction $R-1 \vee R-2 \vee \dots$, the disjuncts of which are related by deep intrinsic resemblances.

Still, matters get even more complicated. For it turns out that different red things satisfy the predicate “red” in virtue of different properties: red light in virtue of its frequency range, red paint in virtue of the chemical composition of its surface, red-hot objects in virtue of their temperature, etc. (see Mellor 2012: 397). So even a seemingly natural property like REDNESS, with much more going for it than YELLOW OR ANGRY may in fact turn out to be almost as disjoint as the latter. If the real existence of properties turns on similarities between disjuncts, then REDNESS may not be a real property at all.

In this respect, what goes for colour-properties also goes for moral properties, which, by virtue of being vague, turn out to be range-properties and therefore surprisingly disjunctive. Could they, unlike colour-properties, be shown to be real in virtue of some deep, intrinsic resemblances between the fine-grained properties constituting their disjuncts? Hardly. Consider once more the property of IMPERMISSIBILITY. It is obvious that actions can satisfy the predicate “impermissible” by virtue of very different underlying properties: some because they are acts of promise-breaking, others because they cause physical harm, still others because they are deceitful, etc., etc. Even thick moral predicates, which are considerably more specific, still manifest the same lack of unity: people can be *courageous* by showing temerity on the battlefield, by coping well through personal tragedies, by braving loss, by making hard decisions, by breaking with old habits, etc.; or they can be *kind* by donating money, by giving up their time, by offering a comforting word, etc. The diversity of properties falling under one and the same moral predicate, be it thick or thin, seems almost boundless. If sameness of property is indeed required, then the revelation that moral properties are multiply realizable by virtue of being vague should cause the realist to question their existence.

But at this stage, the realist has an easy rejoinder at hand: despite their seemingly gerrymandered nature, moral properties could perhaps be acknowledged as real based on a different ontological criterion: not their structure, but their causal efficaciousness. In other words, we should commit ourselves to moral properties if they can pull their weight in explanations and predictions of moral judgments and conduct. The trouble with this move, of course, is that it's only available to naturalists like Jackson (1998) or the Cornell-realists, who have long argued that moral properties can be understood as clusters of more basic, natural properties, and that we should be realists about them so long as they pull their weight in scientific explanations. However, for the non-naturalist this line is totally unpromising.

In addition, there is also the issue of *reduction*, which seems to lead towards a similar conclusion. Given that, on the ontic account, vague moral properties turn out to be co-extensive with disjunctions of sharp properties, could this be used to reduce the former to the latter? The answer depends on what one means by “reduction”, and philosophers have famously disagreed here. If reduction is taken to involve the process of formulating a priori identity-statements that would establish intensional equivalences between the two types of properties, then it should be quite clear that vague properties cannot be reduced to sharp ones. The meaning of “blue” isn’t captured by the disjunction “*B-1* or *B-2* or...*B-17* or...” and it took a piece of a posteriori discovery for us to know what kinds of physical properties underwrite our ascriptions of colours. On the other hand, philosophers like Jackson (1998) or Gibbard (2006), who defend a weaker, non-analytical form of naturalism, could perfectly well argue that, if vague properties can be broken down into disjunctions of sharp properties, then the former are thereby shown to reduce to the latter by virtue of the a posteriori extensional equivalences established. This throws doubt on the claim of Non-reductivism, which is central to non-naturalism.

Such considerations therefore appear to establish that the ontic view of vague properties, while friendly towards various forms of naturalism, remains fundamentally incompatible with non-naturalism about moral properties. There is, however, one final move left for the non-naturalist here—and a rather effective one, too.²¹ On the ontic account of vagueness under consideration, vague properties turn out to be disjunctions of more basic, perfectly sharp properties. But what reason do we have for assuming that the fine-grained properties constituting the disjuncts of such disjunctions are natural rather than non-natural? No argument has been given so far to support this assumption. It is true, of course, that in all the non-moral examples considered (colour-properties, predicates like “tall,” “bald,” etc.) the underlying sharp properties are perfectly natural (wavelengths, heights, numbers and arrangements of hair on people’s heads, etc.). But that doesn’t necessarily mean that the same must obtain in the moral cases. The revelation that the structure of vague moral properties is disjunctive is still compatible with the claim that the fine-grained, sharp disjuncts are themselves irreducibly moral, non-natural properties.

The suggestion that moral ontology may bottom out with a set of perfectly sharp, non-natural properties effectively moves our discussion on to Horn 2 of the dilemma, which starts precisely by assuming that moral properties are sharp. I therefore turn now to exploring this option.

²¹ I owe thanks to David Copp for suggesting this move to me.

3.2. Horn 2: Sharp Moral Properties

The non-naturalist gets into Horn 2 of my dilemma by opting for an account of moral vagueness built around the assumption that moral properties themselves are always sharp. But there are different accounts to be given of sharp properties, which will vary according to which theory of vagueness they are based upon. We have already encountered one such account, derived from the ontic view of vagueness examined in the previous section. I shall start by exploring this account in more detail. Another view of sharp properties, and by far the most popular, is based on a semantic account of predicate vagueness along the lines of supervaluationism: I explore this in §3.2.2. Finally, the third route to sharp properties is one starting from an epistemic account of moral vagueness, the details of which I examine in §3.2.3.

3.2.1. *A Pointilist Moral Ontology?*

Twice during our consideration of Horn 1 above, an idea has emerged according to which the basic constituents of moral reality could be a set of fine-grained, perfectly sharp, irreducibly moral properties. Thus, in §3.1.1 we examined Shafer-Landau's claims about the relation between natural and non-natural properties. According to Shafer-Landau, instantiations of moral properties are constituted by instantiations of natural properties, but the moral properties themselves are not identical to the natural properties that go into their constitution. Following Ridge, I interpreted Shafer-Landau's claims about property-instantiations as claims about tropes. Our second encounter with sharp moral properties came right at the very end of §3.1.2, when we considered an account of vague properties as disjunctive properties. We then asked, on behalf of the non-naturalist: What if the fine-grained disjuncts of such disjunctions were sharp, non-natural properties?

One important insight we have gained from our discussion of Rosen and Smith's account of sharp properties in the previous section is this: for a property to be perfectly sharp, it must be either (i) a maximally specific, point-property, or else (ii) a disjunctive property covering a precisely bound range of such point-properties. Either way, it would seem that, on this picture, the ultimate building blocks of moral reality would be putatively non-natural, maximally specific moral properties. It is, again, tempting to think of these as tropes: concrete, unrepeatable particulars, rather than abstract, universal properties. But these pointilist moral tropes would seem to be strange beasts indeed. It is difficult to even characterize them roughly, for our moral language obviously does not contain terms for qualities so finely grained. Perhaps some approximations would be: "courageous exactly

like that person at *ti*” or “morally wrong precisely like my action at *tj*.”²² As a general strategy, we might try using Kaplan’s “dthat” operator, to yield something along the lines of “dthat [token of *M*]” where *M* would be a term for the kind of property to which the particular trope belongs (“courage,” “kindness,” etc.). But the sense in which such tropes can be said to *belong* to general kinds, as well as the criteria for grouping them together, remains unclear. One obvious option, as we have seen, is to appeal to a relationship of *resemblance* or *similarity* between tropes, based on which we could classify them as belonging to the same kind. The non-naturalist might then argue that, along with moral tropes—to which we must have access via direct acquaintance or some form of intuition²³—our moral experience also delivers basic facts about resemblance. So, for instance, when witnessing Carrie’s act of saving a baby from drowning at great risk to herself, my experience delivers not just the particular moral characteristics of the situation (the moral tropes), but also an objective resemblance to other acts I witnessed in the past (e.g. Dana defending a friend against aggressors), based on which I classify all such acts as instances of “courage.”

But the plausibility of this proposal dissolves under scrutiny. As we saw, resemblances between moral tropes belonging to the same kind (e.g. courage-tropes) could not be resemblances between the natural tropes constituting them. For, just like “red” in Mellor’s example, “courageous,” “cruel,” “kind,” “right,” “good,” etc. can be instantiated by different people or actions in virtue of a myriad of different natural properties. But if moral-trope-resemblance isn’t natural-trope-resemblance, what can it be? The only other option is to appeal to a primitive notion of qualitative resemblance between the moral tropes themselves, over and above any differences and similarities between the natural tropes constituting them. But notice that resemblance is never a simple, all-or-nothing affair: *a* could resemble *b* in some respects but not in others. When two things are similar, it seems natural to assume that they are so because they have certain qualities in common. But once we introduce this idea, the non-realist’s appeal to primitive tropes starts to unravel: how are we to even begin to spell out the putatively non-natural qualities grounding the resemblance between tropes like those exemplified by Carrie’s saving a baby and Dana’s defending a friend, other than by saying that they are both instances of courage? In fact, the very attempt to individuate them qua moral tropes seems to require an appeal to the general kind they belong to: “*courageous* exactly like Carrie’s action at *tj*,” “*courageous* exactly like Dana’s action at *ti*.” But in that case, it turns out after all that the general

²² I thank David Copp for this suggestion.

²³ I assume that since tropes are perfectly simple and unrepeatable, they could not be known via inference, either deductive or inductive.

property COURAGE is primary to, and more basic than, courage-tropes.²⁴ Instead of a pointilist ontology of maximally specific tropes, we're right back where we started: with an ontology of general properties applying across a wide range of cases, and which display all the signs of vagueness. We are back, that is, on Horn 1 of the dilemma. It seems that the non-naturalist must look elsewhere for a satisfactory account of sharp properties.

3.2.2. *Supervaluationism and Sharp Moral Properties*

Thus we arrive at supervaluationism—arguably, the most popular account of vagueness.²⁵ Take a standard vague predicate like “bald,” and let Harry be a borderline instance of that predicate. According to supervaluationism, Harry is a borderline instance of the predicate “bald” because there is no single property that “bald” denotes: rather, there are indefinitely many sharp properties, corresponding to different precisifications of the predicate. Vagueness, on this account, is a matter of semantic indecision: it is indeterminate which of these determinate properties the predicate picks out. Thus, the supervaluationist's world contains, at bottom, only sharp properties.²⁶

It is easy to see, however, that supervaluationism does not hold much promise for the non-naturalist, for a number of reasons. First, there is no room in classic accounts of supervaluationism, such as those developed by Fine (1978) and Keefe (2000), for any reference to properties themselves as distinct from the extensions of predicates. That is to say, the supervaluationist semantics treats properties purely extensionally, as the sets of individuals to which predicates apply. Clearly, this is not the kind of view of properties that will sit well with the non-naturalist's more robust ambitions (and in particular, with the tenet of Non-reductivism).

²⁴ There is, of course, far more to be said here, on both sides of the argument. Dismissing trope-theory isn't something one can do in a quick move like this. But I hope I've done enough to at least raise some doubts about the use of this theory in tandem with moral non-naturalism to yield what I have called a “pointilist moral ontology.” The non-naturalist may have more to say about the relations between moral tropes and properties, and could perhaps use recent work on determinates/determinables, such as Gillett and Rives 2005, to articulate a more robust view. Until such work is completed, however, the doubts I raise are, I think, justified.

²⁵ Thanks to Antti Kauppinen, Tom Dougherty, and an anonymous referee for saving me from a few significant errors I had made in my discussion of supervaluationism in an earlier draft.

²⁶ For a formal explanation of why properties cannot be vague on a supervaluationist account, see Williamson 2003: §5. Keefe (2000: 160) disagrees, but not because she thinks anything significant hangs on whether supervaluationists take properties to be vague or sharp. Quite to the contrary, Keefe believes that, on the extensional view of properties embraced by supervaluationists, it makes no difference whether one says that a predicate “P” indeterminately picks out a sharp property, or rather that “P” determinately picks out a vague property.

But suppose this worry could be alleviated, perhaps by grafting some intensional view of properties onto the supervenience semantics. Would this make supervenience available to the moral non-naturalist? The answer, I think, is still “No.” Assuming that supervenience could be wedded to a robust conception of moral properties, the result would be a luxuriant metaphysics, countenancing a plurality of ever-so-slightly different moral properties, all perfectly sharp and non-natural. Returning to my earlier example, suppose Abe’s action *A* (spending two hours with his patient Ben) is indeterminately permissible. On the proposed picture, we would have to explain this by saying that there are many sharp, ever-so-slightly different non-natural properties of PERMISSIBILITY, some of which include and some of which don’t include *A* as an instance, and it is indeterminate which of those properties the predicate “permissible” picks out. Now, the challenge is for the non-naturalist to explain the odd nature of the moral reality depicted here. This seems difficult.

The non-cognitivist, by contrast, can appeal to a familiar strategy. She will first point out that, on the standard supervenience account (without robust properties added), the source of vagueness is semantic indecision. But of course, as Keefe notes, this isn’t merely “a lazy indecision which is waiting to be resolved” (2000: 155–6). No one thinks we can resolve instances of vagueness, and dissolve all disagreements, simply by fiat. It’s not just that some cases (among which the moral ones are paramount) will resist all kinds of arbitrary stipulation. It’s rather that the very attempt to impose such stipulation is inappropriate and often betrays conceptual confusion. The question, then, is: Why do we find it so deeply inappropriate to resolve borderline cases by fiat? Here, the non-cognitivist can appeal to her favorite explanation, invoking attitudes: as I noted in the Introduction, vagueness is intimately linked with ambivalence. There are many situations in our practical lives when, faced with conflicts of tastes, interests, duties, or even cultures, or perhaps simply with very difficult decisions, ambivalence seems like the right kind of attitude, at least for a while. We do, of course, praise people for showing resolve and the ability to react in the face of adversity. But we also criticize those who are insensitive to the force of the conflict, too opinionated, simple-minded, or plain reckless. We do so because we can hear the force of all the conflicting considerations, and recognize that they are all legitimate. We thus have the rudiments of a psychological account of ambivalence, which could offer someone like Blackburn or Gibbard the starting point in a quasi-realist explanation of moral vagueness.²⁷ On

²⁷ The story, of course, will have to be much more elaborate. But the main ingredients exist: Blackburn does appeal to this kind of ambivalence or open-mindedness (leaving the door open for more than one set of values, or ways of life) in his 1984: 201, though not directly in relation to vagueness.

this quasi-realist account, the higgledy-piggledy nature of moral properties would be taken as a mere reflection of the fact that those properties are projections of our attitudes onto the world.

The non-naturalist, by contrast, has precious little to say about the unruly structure of moral properties like those characterized above. Unlike the quasi-realist, who takes the Euthyphronic direction and explains moral properties in terms of attitudes, the non-naturalist will have to adopt a top-down, Socratic account. But it is unclear to me what the elements of such an account might be.

3.2.3. *Epistemicism and Sharp Moral Properties*

Consider, finally, epistemicism about vagueness.²⁸ Like supervaluationism, this view represents the world as completely precise: there are only sharp properties in the epistemicist's ontology. But instead of viewing vagueness as semantic indecision, this approach characterizes it rather as an epistemic phenomenon: a matter of ignorance concerning the sharp boundaries of our concepts. On this picture, our moral concepts are perfectly sharp: they draw precise boundaries between objects to which they apply (their extensions) and objects to which they don't (their anti-extensions), but it is impossible to know where these boundaries lie due to certain constraints on what counts as knowledge (namely, the "margin-for-error principle").

Take our earlier example of Abe and Ben, where we supposed that Abe's action *A* (spending two hours with Ben) was indeterminately permissible, i.e. a borderline instance of PERMISSIBILITY. What, according to the epistemicist, explains our classifying certain instances of PERMISSIBILITY as "indeterminate" or "borderline," if PERMISSIBILITY itself, as a property, is perfectly sharp? The answer is: ignorance. Although there is a perfectly sharp dividing line between actions that are permissible and actions that are not, we cannot know where that line falls. A nanosecond is enough to make the difference between its being permissible for Abe to spend time with Ben and its not being permissible. Now suppose that *S* and *S** are two stages or time-slices of the universe separated only by one nanosecond, and *p* is the statement that Abe's action *A* is permissible. Suppose, moreover, that the precise boundary falls between *S* and *S**, such that *p* is true at *S* but

²⁸ Shafer-Landau explicitly rejects epistemicism as a plausible explanation of moral vagueness in his (1994), and opts instead for an ontic account. Nevertheless, the view may look independently plausible to other moral realists. Dworkin (2011), for instance, explains apparent instances of moral indeterminacy in terms of ignorance or uncertainty, and insists that there's always a fact of the matter about what we should do. It seems natural to interpret such claims along the lines of epistemicism. And Tim Williamson, who has been one of the most prominent champions of epistemicism, has indicated (in personal correspondence) that he takes a robustly realist stance on morality. It is therefore worth considering the view's metaethical implications in detail.

false at S^* . Could we ever be in a position to know p at S ? No, because our epistemic powers of discrimination are limited: if we are in S , and S^* differs from S only marginally, then for all we know we are in S^* . It follows that if we were entitled to claim knowledge of p at S , then we would have the same entitlement at S^* . But obviously, at S^* our claim to know p would be false. If our belief in p is to count as knowledge, we must have a sufficiently reliable basis to discriminate between p 's being true and p 's being false. Therefore, we are not in a position to know p at S . More generally, in areas where our capacity to discriminate is limited, knowledge requires a margin for error: cases in which one is in a position to know p must not be too close to cases in which p is false. As Williamson puts it, “[w]hen knowing p requires a margin for error, the cases in which p is known are separated from the cases in which p is false by a buffer zone, a protective belt of cases in which p is true but unknown” (Williamson 2000: 18). For the epistemicist, what we call the “grey,” “indeterminate,” or “borderline” area is nothing but this buffer zone of ignorance.

On this view, then, there are determinate moral facts which are strongly unknowable. To many, this is an unpalatable consequence. That a certain moral theory yields unknowable normative truths or obligations is often taken as grounds for rejecting that theory.²⁹ But non-naturalists, like all realists in general, have no problem countenancing unknowable facts and truths: in fact, it's even customary to define realism about a given domain as the claim that truth in that domain is epistemically unconstrained (i.e. can outstrip knowledge). If this holds in other domains of discourse, why would the notion of unknowable truths be any more problematic in the *moral* domain?³⁰

Thus, more needs to be said for the unknowable-truths objection to stick. Yet, it turns out to be surprisingly tricky to put one's finger on the problem. One option would be to appeal to the *action-guidingness* of morals. The main purpose of morality is to provide a guide for action. But how could unknowable moral facts ever be expected to be action-guiding? We can, of course, act in accordance with the reasons they yield—accidentally as it were, by *guessing* at them correctly. But doing so will be essentially a matter of luck. If morality is to be genuinely action-guiding, then we should be able to act morally not just by accident, but by *following* moral rules (in Kantian terms,

²⁹ Among those who find the idea of unknowable obligations objectionable along such lines are Williams 1981; Sider 1995; Thomson 2008; Kramer 2009. Sorensen defends unknowable obligations in his (1995).

³⁰ Shafer-Landau makes this point repeatedly in his (2005), esp. in ch. 10. On the other hand, in his earlier (1994) he distances himself from the epistemicist's unknowable truths when he says: “As a metaphysical realist, I can countenance some unknowable truths. But the number of such truths required by [epistemicists] is so great as to cast doubt on the plausibility of their view.”

we could say: not merely *in accordance with*, but *from* duty). But acting *from* duty involves being in a certain state of mind: *knowing* the contents of our fact-given obligations. Recently, Dougherty writes in a similar vein:

[O]ne might hold that if it is a fact that you must take a taxi that costs less than \$35.41 in order to keep a promise to meet someone at a particular time, then you must be able to decide to take a \$35.41 taxi on the basis of this fact. That is, arguably, the weakest link one could find between ethical facts and motivation. But friends of epistemic explanations of ethical vagueness should deny that even this link obtains. This is because ethical facts that are unknowable could not guide action. (Dougherty 2013: 10)

The non-naturalist seeking to adopt an epistemic theory of vagueness is not likely to be very impressed by such objections, though. For one thing, the scope of unknowable facts should not be overstated: that there are some unknowable truths within a grey area of radical ignorance doesn't impugn on the whole realm of moral facts outside that area. But even focusing on facts in the grey area itself, the non-naturalist can bite Dougherty's bullet and simply deny any link between moral facts and motivation: after all, if moral facts aren't supposed to be projected, or constructed, from facts about our motivational states, then why expect them to always be responsive to such states? Call Dougherty's taxi-involving moral fact "M." According to Dougherty, one might expect that if *m* obtains, one should be able to decide to act as M requires on the basis of M itself. But why should that be a valid expectation? Presumably, by "deciding on the basis of a fact" Dougherty means "deciding on the basis of one's knowledge of that fact." That, although ideal, is hardly a requirement: often, we are forced to decide solely on the basis of our *beliefs* about the facts. And deciding to act as M requires based on one's *belief* that M is certainly possible in Dougherty's scenario.

The difficulty, I believe, comes from focusing on motivation. One way in which moral facts can be action-guiding is by producing corresponding motives to act. But aside from motivation, action-guidingness can also mean providing a *justification*, and by focusing on justification we can uncover, I believe, a deeper problem with the notion of unknowable moral facts. The problem is, roughly, that it is much harder to drive a wedge between moral facts and justification than it is to separate facts from motivation.

The distinction between justification and motivation corresponds to that between *normative* reasons and *motivating* reasons. Normative reasons are justifying considerations. When we have such reasons, and act for them, they become our motivating reasons.³¹ Call these "n-reasons" and "m-reasons," respectively. Now, according to the thesis of Rationalism, as embraced by Shafer-Landau, Parfit, and others, moral facts are intrinsically

³¹ Here I follow Parfit 2011: i. 37.

reason-giving. What kinds of reasons are meant here? Should we think of moral facts as intrinsically m-reason-giving? We could (some realists do), but needn't. It seems more plausible to interpret the thesis instead in terms of n-reasons: moral facts intrinsically generate n-reasons; an agent has the corresponding m-reasons when the agent appropriately relates to her n-reasons.

Now, return to our Abe and Ben example once more. According to epistemicism, there is a sharp boundary between the actions that are permissible for Abe and those that are not, and the difference can be as minute as one nanosecond. Suppose the line between PERMISSIBLE and IMPERMISSIBLE falls between 1.44×10^{13} and $(1.44 \times 10^{13}) + 1$ ns. That is, it would be permissible for Abe to spend four hours with Ben, but one nanosecond more and he would be acting wrongly. It is therefore a perfectly determinate moral fact, call it "F," that it's wrong for Abe to stay with Ben for $(1.44 \times 10^{13}) + 1$ ns. But Abe cannot know F, due to the margin-for-error principle encountered above: any justification that would entitle Abe to claim knowledge of F would also entitle him to claim knowledge of the fact that it's wrong to stay with Ben for 1.44×10^{13} ns. But the latter isn't a fact. So F is unknowable for Abe. Now, according to Rationalism, F generates an n-reason for Abe to stop his emergency session with Ben after at most 1.44×10^{13} nanoseconds (call this " ϕ "). However, Abe has no justification for believing he ought to ϕ , for he cannot know that the line between PERMISSIBLE and IMPERMISSIBLE falls there. We have reached the following conclusions: there is an F-given n-reason for Abe to ϕ , but no justification for him to ϕ . This seems almost contradictory, since n-reasons are supposed to be justifying considerations in favour of actions. But the contradiction is merely apparent: we can distinguish between "There is an F-given n-reason for Abe to ϕ " and "Abe has an F-given n-reason to ϕ ." In our case, the latter is false but the former can well be true.

It makes sense, of course, to separate one's n-reasons from one's actual epistemic state in this way. But what doesn't seem possible is to divorce n-reasons even from a maximally improved capacity for practical rationality. Thus, we can of course accept that there may be moral reasons for us to desist from some of our current practices, but that those reasons are inaccessible to us, due to certain biases or errors in our judgement of which we are unaware. But to recognize them as reasons means to accept that they would serve as justifications for us *if our reasoning abilities were improved*. What seems incoherent is the thought of an n-reason entirely divorced even from the sound exercise of a *maximally improved* capacity for practical reasoning.³² To claim that there are reasons which couldn't be anyone's reasons

³² I draw quite substantive inspiration here from Lillehammer 2002, which provides one of the most forceful arguments I know for taking reasons to be essentially tied to the exercise of our deliberative capacities.

seems almost vacuous. I shall express this upshot in the form of the following epistemic constraint on normative reasons:

Epistemic constraint on reasons: If R is an n-reason for X to Φ , then R can feature in a rational justification of the claim that X ought to Φ , a justification which X knows or could come to know if X's reasoning abilities were maximally improved.

Non-naturalists can, and often do, acknowledge this conceptual link between reasons and our maximally improved reasoning abilities. Thus, Parfit defends an objectivist view of reasons according to which "when it is true that we have decisive reasons to act in some way, this fact makes it true that if we were fully informed and both procedurally and substantively rational, we would choose to act in this way" (2011: i. 63). Similarly, in the course of defending the notion of unknowable moral facts (though not in the context of vagueness), Shafer-Landau concedes that "realists are *not* committed to the idea that moral truths are inaccessible to absolutely ideal epistemic agents at the Piercean limit of enquiry. Epistemically ideal agents who have reached this limit will be fully informed. This means that they will know all facts. Moral realists believe that some of these facts are moral ones; so a genuinely ideal epistemic judge will know all moral facts" (2003: 17).

Now, on an epistemic account which takes vagueness to be just ignorance concerning the sharp extensions of our moral concepts, this ignorance is *irremediable*: no improvement in our rational abilities could remove this uncertainty, due to the margin-for-error constraints attaching to knowledge. So there are *radically* unknowable moral facts which not even agents with maximally improved rational capacities can come to know. If that is so, then either those facts fail to generate n-reasons, in which case Rationalism is false, or else there are reasons which cannot feature in an intelligible justification for anyone, violating the epistemic constraint on reasons above.

There is a reply here on behalf of the non-naturalist, which relies on questioning the notion of a "maximally improved rationality" that I have been using so far.³³ Why assume that an ideal agent, who took every opportunity of improving and refining her reasoning abilities (both instrumental and substantive) and knows all the facts, would still be ignorant about the boundaries of our moral concepts? In the case of vague concepts whose application depends on perceptual criteria ("red," "tall," "heap," etc.), it makes sense to think that even someone possessed of a perfect capacity for reasoning would still remain ignorant about the relevant sharp boundaries,

³³ Thanks to Jen Hornsby, Michael Garnett, and Sarah McGrath for independently alerting me to this possibility.

because the margins-for-error at play are determined by our imperfect *powers of perceptual discrimination* and there is no reason to think that perfect rationality increases one's ability to perceptually discriminate down to nanometers or nanoseconds. However, the non-naturalist might hold that things are different in the moral realm: if moral properties are non-natural, it might be the case that we have epistemic access to them through some form of rational intuition. But then, to imagine an agent whose rational capacities have maximally improved just is to imagine an agent with perfect moral discrimination. In that case, we could tie our conception of reasons to this image of the ideal agent, and preserve Rationalism by claiming that all facts, including the (currently) unknowable, are reason-giving.

Suppose now that Mia is such an ideal agent. Mia woke up one morning and found that she had perfect rationality and unfailing moral discrimination. She knows now precisely when an action stops being permissible, when a white lie turns into a guilty lie, when a fetus becomes a person, how much kindness is too much, etc. Everyone around her still struggles, as per usual: people have a hard time deciding what to do, when to tell a white lie, how much to give to charity, who needs their help more, etc. Abe, for instance, has just spent almost four hours trying to help Ben overcome his panic attack, but has promised his family he'll be home in the afternoon. He doesn't know how much more he can benefit Ben by being there, nor how Ben's benefit would weigh up against the distress his absence is causing at home. Mia, however, knows that the line between permissibility and impermissibility falls precisely at 1.44×10^{13} nanoseconds. Abe is coming up against it quite quickly, so she knows that he should get ready to go back home. But he has no way of knowing that, and she remembers just how difficult these choices used to be. Can Mia say that there's a *normative reason* for Abe to leave soon, despite the fact that he really has no way of knowing it? That's probably neither here nor there. But suppose now that Abe leaves after $1.44 \times 10^{13} - 1$ nanoseconds, i.e. one nanosecond before the precise boundary between permissibility and impermissibility. Mia is happy and praises Abe. If he'd been one nanosecond late, his action would have been wrong. Mia would have blamed and chastized him.

Something here is amiss. Is it right for Mia to judge Abe so differently in the two situations imagined, despite the fact that in one he leaves just one nanosecond later than in the other? Considerations like this have led some to argue that, in the moral realm, we should reject the epistemicist's idea of precise moral properties *on moral grounds*. In general, it seems that if two people (actions, situations, etc.) X and Y are almost absolutely indistinguishable, with the exception of a minute difference with respect to one of the considerations influencing our moral appraisals, then it would be unjust for X to receive different treatment or to be appraised differently than

Y. But that is exactly what happens in a case of vagueness if the epistemic account is correct: the slightest difference (one nanosecond, one nanogram, one nano-anything) is all it takes for an option to change moral valence. This appears to violate moral considerations pertaining to justice. Based on such considerations, Matthew Kramer (2009) has recently argued against vagueness-epistemicism and its attending concept of unknowable obligations, by claiming that such notions would make morality “ludicrous” and “unfair.” Kramer works within a theoretical framework that collapses any distinction between ethics and metaethics, insisting instead on treating the issues of the latter domain as issues belonging to the former. Thus, the question of whether there are any unknowable moral obligations becomes, for Kramer, a question concerning our ethical commitments: should we countenance such obligations in our moral system? What would be the ethical implications of doing so? Because he thinks unknowable obligations would be unjust, Kramer argues that there can’t be any, and therefore that the moral realist would do well to accept genuine moral vagueness, without attempting to reduce it to uncertainty. But this move is not likely to impress the non-naturalist, who takes the business of moral theorizing to be not that of constructing, but of discovering moral truths. Since we don’t create moral facts, why assume that we are at liberty to countenance or reject some of them from our ontology *on moral grounds*?

I believe, however, that Kramer’s point can be restated in language that is more likely to impress the robust non-naturalist. Consider again Shafer-Landau’s conception of the relation between natural facts and non-natural facts:

According to the sort of ethical non-naturalism that I favour, a moral fact supervenes on a particular concatenation of descriptive facts just because these facts realize the moral property in question. Moral facts necessarily covary with descriptive ones because moral properties are always realized exclusively by descriptive ones. Just as facts about a pencil’s qualities are fixed by facts about its material constitution, or facts about subjective feelings by neurophysiological (and perhaps intentional) ones, moral facts are fixed and constituted by their descriptive constituents. (2003: 77)

Consider Shafer-Landau’s pencil example. The pencil’s qualities supervene upon its material constitution, meaning that there couldn’t be a change in the pencil’s qualities without a change in material constitution. But it’s a relatively overlooked fact that supervenience seems to also obey a quantitative requirement of proportionality. This means that there cannot be a great change in the supervening properties without a great change in the subvening base. Thus, a minute change in one of the pencil’s atoms cannot make it the case that the pencil is now rubbish when before it wrote perfectly fine; nor can the pencil turn from long to short by losing just one nanometer.

I am therefore inclined to think that we should uphold the following constraint on the supervenience relation:

Proportionality constraint on supervenience. If P-properties supervene on Q-properties, then no two things can differ greatly with respect to their P-properties without differing greatly also with respect to their Q-properties.

Armed with this constraint, it now becomes quite evident that Kramer's point can be restated in metaphysical rather than purely moral terms. The fact that it's unjust to praise one person and blame another when the difference between their actions was slight is, of course, a moral consideration. But the consideration is grounded, I believe, in the proportionality constraint on supervenience: moral responsibility, praise, and blame are concepts that supervene on the natural properties instantiated by those people's actions. If the difference in those natural properties was small, so must be the difference in their moral properties, too.

I conclude, based on all of this, that epistemicism about vagueness isn't an option for the non-naturalist seeking to preserve Supervenience. We have three accounts of sharp properties in this section, and found them incompatible with non-naturalism. It looks as if the non-naturalist can appeal neither to theories positing vague properties, nor to accounts based on sharp properties, in his attempt to elucidate moral vagueness.

The dilemma appears to stick.

4. CONCLUSIONS

I have argued in this chapter that the moral non-naturalist seeking to countenance moral vagueness in an attempt to explain moral disagreement faces a dilemma. Non-naturalism I have described as commitment to seven theses: Cognitivism, Correspondence, Atomism, Objectivism, Supervenience, Non-reductivism, and Rationalism. On either horn of the dilemma, serious problems arise for some of these theses: in various ways, vague properties seem to threaten Objectivism, Supervenience, and Non-reductivism; on the other hand, sharp properties raise problems for Supervenience and Rationalism. The difficulties on each horn of the dilemma are real, and while they may not be insuperable, they do, at the very least, drastically limit the things non-naturalists can consistently say about moral properties, facts, and reasons. Non-naturalism may in the end survive my dilemma, but if it does it will be a doctrine considerably different from what some of its leading proponents take it to be.³⁴

³⁴ I am indebted to Maïke Albertzart, Tom Dougherty, Antti Kauppinen, Laura Vasile, and two anonymous reviewers for Oxford University Press for extremely helpful comments on earlier drafts of this chapter. Previous versions of the material were presented

References

- Aristotle. 2009. *Nicomachean Ethics*, tr. D. Ross. Oxford and New York: Oxford University Press.
- Armstrong, D. M. 1978. *A Theory of Universals: Universals and Scientific Realism*, ii. Cambridge and New York: Cambridge University Press.
- Audi, R. 2004. *The Good in the Right: A Theory of Intuition and Intrinsic Value*. Princeton: Princeton University Press.
- Crisp, R. 2006. *Reasons and the Good*. Oxford: Oxford University Press.
- Cuneo, T. 2007. "Recent Faces of Moral Nonnaturalism," *Philosophy Compass*, 2: 850–79.
- Dancy, J. 1981. "On Moral Properties," *Mind*, 90: 367–85.
- Dancy, J. 2006. "Nonnaturalism", in D. Copp (ed.), *Oxford Handbook of Ethical Theory*, 122–45. New York: Oxford University Press.
- Dougherty, T. 2013. "Vague Value," *Philosophy and Phenomenological Research* (published online first, DOI: 10.1111/phpr.12026).
- Dworkin, R. 1977. "No Right Answer?" in P. Hacker and J. Raz (eds), *Law, Morality, and Society*, 58–84. Oxford: Clarendon Press.
- Dworkin, R. 2011. *Justice for Hedgehogs*. Cambridge, MA, and London: Harvard University Press.
- Enoch, D. 2011. *Taking Morality Seriously: A Defense of Robust Realism*. Oxford: Oxford University Press.
- FitzPatrick, W. 2008. "Robust Ethical Realism, Non-Naturalism, and Normativity," in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, iii. 159–205. Oxford: Oxford University Press.
- Gibbard, A. 2006. "Normative Properties," in T. Horgan and M. Timmons (eds), *Metaethics After Moore*, 319–38. Oxford: Clarendon Press.
- Gillett, C., and Rives, B. 2005. "The Non-Existence of Determinables: Or, a World of Absolute Determinates as Default Hypothesis," *Noûs*, 39: 483–504.
- Huemer, M. 2005. *Ethical Intuitionism*. New York: Palgrave Macmillan.
- Hurley, S. 1992. *Natural Reasons: Personality and Polity*. Oxford: Oxford University Press.
- Hyde, D. 2008. *Vagueness, Logic and Ontology*. Aldershot: Ashgate.
- Jackson, F. 1998. *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Clarendon Press.

at the 2012 Madison Metaethics Workshop, and at a departmental seminar at Birkbeck College, University of London. I am grateful to audiences at both events, and in particular to David Copp, Terence Cuneo, James Dreier, Dorothy Edgington, Michael Garnett, Jen Hornsby, Keith Hossack, Sarah McGrath, Robert Northcott, Anthony Price, Russ Shafer-Landau, and Robert Shanklin. Hallvard Lillehammer read a very early ancestor of this chapter many years ago, and gave me lots of insightful suggestions which provided the impetus for much of the work done here.

- Keefe, R. 2000. *Theories of Vagueness*. Cambridge: Cambridge University Press.
- Kramer, M. 2009. *Moral Realism as a Moral Doctrine*. Oxford: Wiley-Blackwell.
- Lewis, D. 1986a. *On the Plurality of Worlds*. Oxford: Basil Blackwell.
- Lewis, D. 1986b. "Against Structural Universals," *Australasian Journal of Philosophy*, 64: 25–46.
- Lillehammer, H. 2002. "Moral Realism, Normative Reasons, and Rational Intelligibility," *Erkenntnis*, 57: 47–69.
- Mellor, D. H. 2012. "Nature's Joints: A Realistic Defence of Natural Properties," *Ratio*, 25: 387–404.
- Moore, G. E. 1903. *Principia Ethica*. Cambridge: Cambridge University Press.
- Parfit, D. 2011. *On What Matters*. Oxford: Oxford University Press.
- Plato. 2000. *The Republic*, ed. G. R. F. Ferrari, tr. T. Griffith. Cambridge: Cambridge University Press.
- Railton, P. 1992. "Pluralism, Determinacy, and Dilemma," *Ethics*, 102: 720–42.
- Rosen, G., and Smith, N. J. J. 2004. "Worldly Indeterminacy: A Rough Guide," *Australasian Journal of Philosophy*, 82: 185–98.
- Russell, B. 1923. "Vagueness," *Australasian Journal of Philosophy and Psychology*, 1: 84–92.
- Sainsbury, R. M. 1989. "What is a Vague Object?," *Analysis*, 49: 99–103.
- Sainsbury, R. M. 1996. "Concepts without Boundaries," in R. Keefe and P. Smith (eds), *Vagueness: A Reader*, 251–64. Cambridge, MA: MIT Press.
- Sanford, D. H. 1966. "Red, Green, and Absolute Determinacy," *Philosophical Quarterly*, 65: 356–8.
- Sanford, D. H. 1970. "Disjunctive Predicates," *American Philosophical Quarterly*, 7: 162–70.
- Sanford, D. H. 2013. "Determinates vs. Determinables," in E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. (spring 2013 edn): <<http://plato.stanford.edu/archives/spr2013/entries/determinate-determinables>>.
- Scanlon, T. M. 2003. "Metaphysics and Morals," *Proceedings and Addresses of the American Philosophical Association*, 77: 7–22.
- Schiffer, S. 2000. "Vagueness and Partial Belief," in E. Sosa (ed.), *Philosophical Issues*, x. *Skepticism*, 220–57. Oxford: Blackwell.
- Schiffer, S. 2010. "Vague Properties," in R. Dietz and S. Moruzzi (eds), *Cuts and Clouds: Essays on the Nature and Logic of Vagueness*, 109–30. Oxford: Oxford University Press.
- Shafer-Landau, R. 1994. "Ethical Disagreement, Ethical Objectivism, and Moral Indeterminacy," *Philosophy and Phenomenological Research*, 54: 331–44.
- Shafer-Landau, R. 1995. "Vagueness, Borderline Cases and Moral Realism," *American Philosophical Quarterly*, 32: 83–96.
- Shafer-Landau, R. 2003. *Moral Realism: A Defence*. Oxford: Clarendon Press.
- Shapiro, S. 2006. *Vagueness in Context*. Oxford: Clarendon Press.
- Sider, T. 1995. "Sorensen on Unknowable Obligations," *Utilitas*, 7: 273–9.
- Sorensen, R. 1990. "Vagueness Implies Cognitivism," *American Philosophical Quarterly*, 27: 1–14.
- Sorensen, R. 1995. "Unknowable Obligations," *Utilitas*, 7: 247–71.

- Sosa, E. 2001. "Objectivity without Absolutes," in A. Byrne, R. C. Stalnaker, and R. Wedgwood (eds), *Fact and Value: Essays on Ethics and Metaphysics for Judith Jarvis Thomson*, 215–27. Cambridge, MA, and London: MIT.
- Thomson, J. J. 2008. *Normativity*. Chicago: Open Court Press.
- Tye, M. 1990. "Vague Objects," *Mind*, 99: 535–57.
- Vasile, L. J. 2010. "Moral Disagreement and Moral Indeterminacy," unpublished doctoral dissertation, Department of Philosophy, Central European University.
- Wedgwood, R. 2007. *The Nature of Normativity*. Oxford: Oxford University Press.
- Williams, B. 1981. *Moral Luck*. Cambridge: Cambridge University Press.
- Williamson, T. 1994. *Vagueness*. London: Routledge.
- Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.
- Williamson, T. 2003. "Vagueness in Reality", in M. J. Loux and D. Zimmerman (eds), *Oxford Handbook of Metaphysics*, 690–715. Oxford: Oxford University Press.
- Wolf, S. 1992. "Two Levels of Pluralism," *Ethics*, 102: 785–98.