

# Ethical Vagueness and Practical Reasoning\*

Billy Dunaway  
University of Missouri–St Louis

Draft of 23 October 2015  
Forthcoming in *The Philosophical Quarterly*

## Abstract

This paper looks at the phenomenon of ethical vagueness by asking the question, how ought one to reason about what to do when confronted with a case of ethical vagueness? I begin by arguing that we must confront this question, since ethical vagueness is inescapable. I then outline one attractive answer to the question: we ought to maximize expected moral value when confronted with ethical vagueness. This idea yields determinate results for what one rationally ought to do in cases of ethical vagueness. But what it recommends is dependent on which substantive theory of vagueness is true; one can't draw conclusions about how to reason about vagueness in ethics in the absence of concrete assumptions about the nature of vagueness.

Vagueness—or indeterminacy, or indefiniteness—is pervasive, and the practical domain is not exempt from it. Cases where it is vague what we *ethically* ought to do are widespread. One response to this situation is to acquiesce in ethical vagueness and let related notions of practical normativity, blameworthiness, rationality, go vague whenever vagueness in ethics is present. On this approach, when there is vagueness in the ethical status of an action, it immediately follows that it is also vague whether one rationally ought to perform that action, it is vague whether one is blameworthy if one performs it, and so on. But this isn't the only approach to the issue, and I will set it aside for present purposes. After giving some structure to the issue of ethical vagueness, I outline an approach to achieving definite answers to the question of what one rationally ought to do, when it is vague what one ethically ought to do.

---

\*Special thanks to Elizabeth Barnes, Ross Cameron, Jennifer Carr, Matti Eklund, John Hawthorne, Miriam Schoenfield, Mark Schroeder, Alex Silk, Julia Staffel, David Manley, Robbie Williams, and participants at the Rocky Mountain Ethics Congress and the Indeterminacy in What We Care About conference at the University of Leeds for discussion of the issues in this paper.

## 1 Cases of ethical vagueness

We cannot avoid these questions for practical reasoning posed by ethical vagueness by denying that ethical vagueness exists. One hallmark of vague terms is susceptibility to a sorites series. And sorites series for ethical terms are easy to find.

### 1.1 The general argument

It is permissible to kill an amoeba to save a modern, healthy, and fully functioning adult member of the species *homo sapiens*. And it is wrong to kill another equally healthy human to save the first human. But there will be some creatures from along the line of human evolutionary history for which it is implausible to say that it is wrong to kill the creature in such a case, and also implausible to say that it is not wrong to kill the creature in such a case. These will be cases where it is vague whether killing the creature is wrong.<sup>1</sup>

In more detail: let  $c_1 \dots c_n$  be a series of cases where in each there is a creature in front of you, and a button you could press which would kill the creature to save the healthy adult human being. If you don't press the button and kill the creature, the human dies. The creatures in each case are all evolutionary ancestors of human beings. They are arranged in order of ascending complexity: in  $c_1$ , it is a pre-historic amoeba in front of you, which you can kill to save the human, and in  $c_n$ , another human is in front of you. And for each case between  $c_1$  and  $c_n$ , another creature from human evolutionary history is in front of you, slightly more developed than the creature in the case before it. (That is, in  $c_2$  a creature from our evolutionary history slightly more developed than the amoeba is there; in  $c_3$  it is a slightly more developed creature still, and so on, until in  $c_{n-1}$  an immediate ancestor of modern humans is in front of you.) Thus, for each  $c_i$  and  $c_{i+1}$  ( $0 < i \leq n$ ), the creature in front of you in  $c_{i+1}$  is slightly more developed than the creature in front of you in  $c_i$ .

This constitutes a sorites series for 'wrong', since the following three claims hold. First,  $c_1$  is a case where it is clearly permissible to press the button. Second,  $c_n$  is a case where it is clearly wrong to press the button. And third, for any case  $c_i$  in between  $c_1$  and  $c_n$ , there is significant pressure not to assert the conjunction of the following claims:

**Wrong<sub>*i*</sub>** It is not wrong to press the button in  $c_i$ ;

**Wrong<sub>*i+1*</sub>** It is wrong to press the button in  $c_{i+1}$ .

Thus 'wrong' is similar to paradigmatically vague terms like 'bald', where similar claims apply. First, someone with no hairs on their head is bald. Second, someone with 1,000,000 hairs on their head is not bald. And third, for any number

---

<sup>1</sup>For more on sorites-series as a theory-neutral characterization of vagueness, see Greenough (2003).

$i$  between 0 and 1,000,000, there is significant pressure not to assert the conjunction of the following claims:

**Bald <sub>$i$</sub>**  A person with  $i$  hairs is bald;

**Bald <sub>$i+1$</sub>**  A person with  $i + 1$  hairs is not bald.

This is strong evidence that 'wrong' is vague.<sup>2</sup>

It is useful to think about how series of cases with this structure arise in matters close to everyday life. One can think of the development of a sperm and egg immediately prior to fertilization, which then becomes a human child, in similar terms. It is permissible to destroy the sperm and egg prior to fertilization; it is impermissible to destroy the child. But there will be some points along the development of the organism where it is vague whether it is wrong to destroy the thing that will become the child. Similarly with self-regarding duties and enhancement: drinking a cup of coffee is permissible; taking a drug that gives one the same cognitive boost but then kills you in a month isn't. Passive euthanasia provides another case: not undertaking a small task to prolong the life of a person with many years of life ahead of them is wrong; not undertaking the same task to prolong for one second the life of someone with a painful disease is permissible. Examples of this kind abound; the purpose of the present paper isn't to enumerate them, but it bears keeping in mind that the conclusions of this paper have straightforward application to many everyday examples of this kind.

## 1.2 *Precise measurement and absolute vs. comparative ethical predicates*

One might be tempted to think that the above sorites series for ethical predicates depends on some specific assumptions about ethical predicates that are in principle dispensable. In particular, one might suspect that it depends on the assumption that ethically relevant properties cannot be precisely measured and compared. Or, one might suspect that it depends on the assumption that ethical predicates are not comparative in structure. I will briefly sketch below why these suspicions would be unwarranted.

First, the existence of ethical vagueness has little to do with the unavailability of precise measurements of ethically relevant properties, or of incomparability among such properties. Assume for the moment that a simple Utilitarian theory in the style of Bentham (1781) is true, on which the only ethically relevant property of an action is how much net utility it produces; the right action is, on this theory, the one that produces the greatest amount of net utility. Assume that only happiness contributes to positive utility, and pain to negative utility. And assume that quantities of happiness and pain can be precisely measured (so that if an act produces some utility, it can be assigned a real number corresponding to the amount of positive or negative utility produced) and compared (so that one act produces more net utility than another iff the real number assigned to the

---

<sup>2</sup>See Shafer-Landau (1995) for a similar conclusion, though sorites series are not his focus.

former is greater than the latter). Even given all of these assumptions, there will be ethical vagueness. This is because it can be vague whether an act produces any positive or negative utility at all.

Here is a sorites series illustrating this: imagine a series of cases, in which, much like before, one can press a button to save a human. But let the button in each case operate as follows: in  $c_1$ , it destroys 1,000,000 amoebas, in  $c_2$ , it destroys the same number of slightly more complex creatures, and in  $c_n$  it destroys 1,000,000 humans who are intrinsically the same as the human to be saved. On the version of Utilitarianism spelled out above, it is permissible to press the button in  $c_1$  since no negative utility is produced by killing an amoeba, and significant positive utility is gained by saving a human. And it is wrong to press the button in  $c_n$ . But there will be cases in between  $c_1$  and  $c_n$  where it is vague whether the creatures that will be destroyed by the button can experience happiness or pain at all. In these cases it is vague how much net utility pressing the button produces (and vague how much net utility refraining from pressing the button produces) hence it is vague, according to the view under consideration, whether it is wrong to press the button in these cases.<sup>3</sup>

It is also worth mentioning that vagueness in ethics does not depend on the assumption that ethical predicates are non-degreed predicates which serve to pick out a threshold on a scale of ethically relevant factors.<sup>4</sup> To be sure, this is one way in which ethical vagueness might arise: if 'wrong' applies to those actions that meet some threshold (for instance, if some amount of autonomy violation constitutes wrongness), then it will be very natural to think that there is no precise threshold at which 'wrong' begins to apply.<sup>5</sup> But it would be a mistake to think that this is the only source of ethical vagueness. Even for comparative or degreed predicates ('A is more wrong than B'; 'X is wrong to degree n') there will be sorites series beginning with a case to which the relational predicate applies and ending with a case where it clearly does not apply.<sup>6</sup> Vagueness cannot be escaped by going gradational.

## 2 What to do when it's vague: definitions and assumptions

I will be exploring possible answers to the question, 'what ought one to do when it is vague whether an available action is wrong?'. I will be supposing that we are trying to answer this question in a context where ethical considerations are the only considerations that bear on performing the action. In other words, when discussing cases where it is vague whether an action is ethically wrong, it will be a background assumption that in the cases in question any prudential or non-

---

<sup>3</sup>Compare Shafer-Landau (1995).

<sup>4</sup>See Scanlon (1998) and Schroeder (2007) for gradational approaches to ethical facts in terms of weighted reasons. Thanks to Mark Schroeder and an anonymous referee for discussion of this issue.

<sup>5</sup>Dougherty (2013: 2)

<sup>6</sup>See Keefe (2000: 12-15) for a more extended argument.

ethical reasons for performing available options are not relevant. Either they do not exist, or the reasons in favor of performing an action are balanced by the reasons against. Thus there is no possibility of answering the question of whether I should perform a vaguely wrong action in the negative, on the grounds that it will make me feel slightly cold.

Here then is one way to think about the question ‘what ought one to do when it is vague whether an available action is wrong?’. ‘Ought’ can take on various senses, one of which is ethical. But there are other senses as well, and vagueness in the ethical sense need not imply vagueness in other senses. Among these other senses is an ‘ought’ that picks out the actions that are best in view of what one knows.<sup>7</sup> More specifically, I will be understanding this ‘ought’ as follows: ‘ought  $\phi$ ’ is true just in case  $\phi$ -ing is best, in view of what one knows. (If multiple actions tie for best, then each is permissible.) This ordering on actions is naturally heard as the one at issue on the true reading of the sentence ‘Sally ought to bet on heads’ when she is offered a bet with identical payouts on outcomes of a coin flip with a coin she knows has a bias  $\beta$  (where  $\beta > .5$ ) in favor of heads. Since it is more likely on what Sally knows that the coin will land heads and the payoffs of winning on a ‘heads’ bet are the same as the payoffs of winning on a ‘tails’ bet, the action of betting ‘tails’ isn’t among the best actions in view of what she knows. Let us call this the “rational ‘ought’”, though this may be misleading as there are likely other uses of ‘rational’ that don’t line up with the one I outline here.

We can then ask, what ought rationally an agent to do when it is vague whether an action available to her is wrong? With the assumptions outlined above, this amounts to the question: when faced with an action that is vaguely wrong, which actions are best, in view of what one knows? Focusing on the sorites series from §1, this is the question of whether pressing the button is best in view of what one knows, when it is vague whether pressing the button is wrong.

The following background assumptions about what is known in these cases, which are independent of any theory of vagueness, will be in place throughout:

First, we will assume that one is certain about which ethical theory is correct. The existence of a sorites series of the kind described in §1 does not depend on uncertainty in this matter: even if one were certain that one ethically ought to maximize net utility, there are still possible cases where one would find it vague whether one ethically ought to press the button. Of course if we treat every possible precisification of ‘maximizes net utility’ as a distinct ethical theory, one can’t necessarily know which of *these* theories is correct. Vagueness might preclude knowing that it is wrong to do an action with the property instantiated in case  $c_{i+1}$  but not  $c_i$  in a sorites series (more on this in §3). But at least at the level at which ethical theories are traditionally specified—e.g., ‘Utilitarianism’, ‘Kantian deontology’, etc.—I will assume that the correct theory is known.<sup>8</sup>

---

<sup>7</sup>For more, see Kratzer (1977).

<sup>8</sup>Hence we are not in a case of ethical uncertainty of the kind discussed by Ross (2006) and Sepielli (2009).

The second assumption is that one knows all of the physical facts about each case in a sorites series. That is, one knows the microphysical differences between each case in the series, what will result (microphysically) if one performs available action in a case, and the like. There very well could be some vagueness in the microphysical facts too, and in which microphysical fact “constitutes” a case, but I will ignore them for convenience. Instead, I will focus on the following: even someone who was fully able to comprehend the microphysical nature of a case, and how the world would unfold microphysically if a particular action were to be performed, would still find it vague in some cases whether a particular action produces happiness, violates autonomy, etc.

The third assumption is that one knows certain necessary truths about each borderline case. For each borderline case  $b_i$  in a sorites series for ‘wrong’ where it is vaguely wrong to press the button in  $b_i$ , there is a wrongness-like property  $Wrong_i$  which has its cut-off point exactly at case  $b_i$ . That is, the property  $Wrong_i$  is such that pressing the button in every case in the series leading up to and including  $b_i$  lacks the property, and pressing the button in every case in the series after  $b_i$  has the property. (For simplicity assume that in other cases,  $Wrong_i$  is instantiated by everything that is wrong, and isn’t instantiated by everything that isn’t wrong.) Letting  $b_1, b_2 \dots b_n$  be the borderline cases in a sorites series, then, the following are *necessary* truths:

WRONG<sub>0</sub>  $b_1, b_2, b_3, b_4 \dots b_n$  are  $Wrong_0$ ;

WRONG<sub>1</sub>  $b_1$  is not  $Wrong_1$  and  $b_2, b_3, b_4 \dots b_n$  are  $Wrong_1$ ;

WRONG<sub>2</sub>  $b_1, b_2$  are not  $Wrong_2$  and  $b_3, b_4 \dots b_n$  are  $Wrong_2$ ;

WRONG <sub>$i$</sub>   $b_1, b_2, \dots b_i$  are not  $Wrong_i$  and  $b_{i+1} \dots b_n$  are  $Wrong_i$ .

In a case of ethical vagueness, one can know the propositions WRONG<sub>0</sub>, WRONG<sub>1</sub>, WRONG<sub>2</sub>, etc. After all, it is presumably a contingent cognitive limitation only that prevents us from using the reference-fixing description used for  $Wrong_i$  above, where the property is referred to by the things that instantiate it. Someone who did go through the cognitive effort needed to refer to properties in this way would know how they apply in borderline cases of wrongness.

The fourth assumption is that the sorites series we are dealing with contain finitely many cases (and, hence, that there are only finitely many borderline cases in a sorites series). Any sorites series for ‘wrong’ doesn’t need infinitely many borderline cases, since the distinctive claim—that for any case there is substantial pressure not to assert that pressing the button in one borderline case is not wrong, but that pressing in the very next case is wrong—will hold so long as the cases in a finitely long series are sufficiently fine-grained. So I will make this assumption to make exposition simpler (i.e., with no need to account for infinite cases) in what follows.

With these assumptions in place, the next section begins by outlining one way of answering the question, what ought one rationally to do when it is vague whether pressing the button is wrong? I begin by explicitly answering the this question while assuming an Epistemicist view of vagueness. I will not be making this assumption not out of a commitment to the plausibility of the Epistemicist view, but rather because it affords an especially simple and workable approach to vagueness which allows for straightforward talk about vagueness which retains classical logic and standard attitudes of belief, knowledge, and credence toward vague cases. It provides the resources to implement a very natural idea: that one rationally ought to maximize expected moral value in borderline cases.

Rational action looks different when we adopt non-Epistemicist approaches to vagueness, but the comparison with Epistemicism is instructive. Some theories must reject that the expected moral value approach provides determinate answers in borderline ethical cases. Others can apply the framework, but deliver different recommendations. Regardless, questions about ethical vagueness and rational action cannot be answered absent some fairly specific theses about what such cases involve.

This is a significant issue, since some important work on practical reasoning and vagueness has treated the question as if it is *independent* of questions about the nature of vagueness. Authors on this topic often do not explicitly rely on substantive views about the nature of vagueness when drawing practical conclusions.<sup>9</sup> The closing sections of this paper point toward the conclusion that no such ecumenical approach to rational action is available.

### 3 What to do when it's vague: maximize expected moral value

#### 3.1 *The framework*

The Epistemicist view in Williamson (1994) is the standard version of the epistemic view of vagueness, and I will for reasons of simplicity begin by approaching the question of what we rationally ought to do in a case of ethical vagueness from within the Epistemicist framework. Take our sorites series for 'wrong' from §1: the Epistemicist view holds that for any case  $c_i$  from  $c_1 \dots c_n$ , there is a fact of the matter whether pressing the button in  $c_i$  is wrong or not. The vagueness in 'wrong' consists in our inability to know, for some of the cases in  $c_i$  from  $c_1 \dots c_n$ , whether pressing the button is wrong or not. These are the borderline cases. Vagueness is, in short, not to be found in language or in the world, but in what we are able to know.

This picture requires some qualification in light of the simplifying assumptions we made above. (We might alternatively try to modify the assumptions, but I

---

<sup>9</sup>See in particular Wasserman (2013: §6), which assumes that we can help ourselves to the Epistemicist's probabilities even if Epistemicism is false, and Williams (2013: 1), which demurs from taking a stand on the "source" of indeterminacy. But see Williams (2014) for a different view, which I discuss in §5.

won't explore this route here.) In particular, we assumed that one can know necessary truths about the properties *Wrong*<sub>1</sub>, *Wrong*<sub>2</sub>, etc. which are defined by where they draw a precise cut-off point in a sorites series. But if 'wrong' itself has a precise cut-off point then wrongness is identical to one of these properties, and so one *can* know where the cut-off point for wrongness is. That is: if the third borderline case, *b*<sub>3</sub>, is the cut-off point for wrongness and *b*<sub>1</sub>, *b*<sub>2</sub>, and *b*<sub>3</sub> are not wrong and *b*<sub>4</sub> . . . are wrong, then the proposition *WRONG*<sub>3</sub> is equivalent to the proposition that *b*<sub>3</sub> is the cut-off point for wrongness. Our assumption that one can know the former is apparently inconsistent with the Epistemicist claim that one can't know the latter.

The inconsistency isn't unresolvable, and in fact there are multiple ways to resolve it. One could deny that the proposition *WRONG*<sub>3</sub> is the same proposition as the proposition that the cut-off point for wrongness is *b*<sub>3</sub>. One can then know *WRONG*<sub>3</sub> but not know that the cut-off point for wrongness is *b*<sub>3</sub> because knowing the first proposition isn't the same as knowing the second. Or, one could hold that these are the same proposition, but entertained under different guises, and that the relevant knowledge attributions are guise-relative. Both approaches have some plausibility, as the property *WRONG*<sub>3</sub> is specified using enumeration of borderline cases, whereas this plausibly isn't how one ordinarily grasps wrongness. I have some sympathy with the guise-theoretic approach, but nothing in the subsequent discussion rests heavily on this assumption.<sup>10</sup>

One important consequence of Epistemicism for practical reasoning in cases of ethical vagueness is the following: while each borderline case is such that we can't know whether pressing the button in it is wrong, each borderline case is not identical to the other borderline cases in every epistemic respect. Instead, for some borderline cases, the likelihood on what one knows that pressing the button in those cases is wrong is much *lower* than the likelihood on what one knows that pressing the button is wrong in other borderline cases. Take for instance the first borderline case *b*<sub>1</sub> and a borderline case further down the sorites series, *b*<sub>*i*</sub>. One can't know that pressing in *b*<sub>1</sub> is wrong, but it is still very unlikely on what one knows that pressing in *b*<sub>1</sub> is wrong. Of all the possible cut-offs for wrongness that one can't rule out, there is just one which includes pressing in *b*<sub>1</sub> among the wrong actions. By contrast, it is much more likely on what one knows that pressing in *b*<sub>*i*</sub> is wrong, since there are more cut-off points which one can't rule out that place pressing the button in this case among the wrong actions. Picturesquely: as one traverses through more borderline cases of wrongness, one should grow more confident that one has passed the cut-off point for wrongness.

This is captured by the following thesis, which we can call *INCREASING CREDENCES*. *Cr*, as I use it below, is a credence function that matches the likelihoods on one's evidence: *Cr*(*p*) is the likelihood of *p* on what one knows. (These are the credences one ought to have, in the sense of the rational 'ought' outlined above.) Then, if *b*<sub>*i*</sub> and *b*<sub>*j*</sub> are borderline cases in a sorites series for 'wrong' where *j* > *i*,

<sup>10</sup>Thanks to John Hawthorne for discussion here.



and  $\mathbf{b}_x$  is the proposition that pressing the button in case  $b_x$  is wrong (under the guise of ‘wrong’), the following is true:

INCREASING CREDENCES  $Cr(\mathbf{b}_j) > Cr(\mathbf{b}_i)$ .

It is very natural to supplement INCREASING CREDENCES with a principle concerning how much one’s credence in each case should increase. For instance: assuming we have a genuine sorites series for ‘wrong’, it would be highly implausible to say given what one knows, the likelihood that pressing the button in the borderline case  $b_i$  is wrong is .3, but that in the very next borderline case,  $b_{i+1}$ , the likelihood that pressing the button is wrong is .7. (If this were rational, it wouldn’t feel so implausible to select  $b_i$  as the point where the cut-off point for wrongness lies, and hence the series wouldn’t be a sorites series.) This amounts to the requirement that one’s confidence that the cut-off point for wrongness lies at a particular borderline case be roughly the same as one’s confidence that the cut-off point is at a distinct borderline case. We can call this the UNIFORMITY thesis which applies to all pairs of borderline cases,  $b_i$  and  $b_j$ :

UNIFORMITY  $Cr(\neg\mathbf{b}_{j-1} \wedge \mathbf{b}_j) \approx Cr(\neg\mathbf{b}_{i-1} \wedge \mathbf{b}_i)$

It is intuitively very plausible that INCREASING CREDENCES and UNIFORMITY are constraints on one’s credences about borderline cases. Borderline cases are structurally similar to other cases where vagueness is not at issue, and in these similar cases analogues of INCREASING CREDENCES and UNIFORMITY are true. Suppose I have drawn a winning ticket from a 10-ticket lottery, where the tickets are numbered sequentially from ‘1’ to ‘10’. Suppose moreover that I have looked at the winning ticket but not told you. You cannot know which ticket won. I then line the tickets up on a table, with the ticket labelled ‘1’ on the left, ‘2’ next to it, and so on, with ‘10’ at the other end. I then start pointing at tickets in order, asking whether a ticket I have pointed at at some time was the winner. You cannot know the answer to this question until I point to ticket 10. But there are some constraints on your credences in this case: let TICKET<sub>1</sub> be the proposition that ticket 1 won, TICKET<sub>2</sub> be the proposition that either ticket 1 or ticket 2 won, and so on. You cannot know any of TICKET<sub>1</sub>-TICKET<sub>9</sub>.<sup>11</sup>

But your credences should have a particular structure: you should be *more confident* in TICKET<sub>2</sub> than you are in TICKET<sub>1</sub>, as the former is more likely on your evidence than the latter. And similarly for any two propositions TICKET<sub>*m*</sub> and TICKET<sub>*n*</sub> where  $10 > m > n$ . Hypotheses according to which one of a large number of tickets was drawn will always be more likely. Moreover for any  $n$  where  $0 < n < 10$ , one’s evidence does not make TICKET<sub>*n+1*</sub> *substantially* more likely than TICKET<sub>*n*</sub>. There is no evidence that any one ticket has a greater chance of being drawn, so adding one more ticket to a hypothesis about which ticket was drawn will increase the likelihood of the hypothesis by roughly the same amount.

---

<sup>11</sup>Cf. Hawthorne (2004)

Hence one's credences in the ticket case should obey analogues of INCREASING CREDENCES and UNIFORMITY.

### 3.2 Moral value and expected value maximization

A case of ethical vagueness induces uncertainty about the ethical facts. And in general In cases of uncertainty, it is very natural to take the rational 'ought' as requiring one to preform the action that maximizes expected value. But the constrains on credences from the previous subsection provide only half of the resources needed by an expected value maximization framework. This subsection sketches how the other half might be filled out, and the recommendations from the rational 'ought' that result.

Let's confine ourselves to a specific ethical theory for the sake of illustration, and a simple explanation of how vagueness arises for this theory. Suppose one is certain that a simple Utilitarianism is correct, and that the vagueness in 'wrong' is evidenced by a sorites series where it is vague whether in some of the cases in the series the amount of net utility produced by pressing the button and killing the creature in front of you is greater than the amount of net utility produced by refraining from pressing the button and letting the human die. According to the Epistemicist, then, one cannot know whether in these cases pressing the button produces more net utility than refraining from pressing it. There are then several aspects to the decision problem of what one rationally ought to do in such a borderline case.

First, in each borderline case  $b_i$ , there are two available actions: pressing the button and not pressing the button. There are also two possibilities given what one knows in each borderline case: that pressing the button would be wrong, and that pressing the button would not be wrong (these propositions are represented as  $\mathbf{b}_i$  and  $\neg\mathbf{b}_i$ , respectively). This gives a total of four possible outcomes, one for each quadrant below:

	$\mathbf{b}_i$	$\neg\mathbf{b}_i$
press		
no press		

Second, given our assumption of a simple Utilitarianism, the amount of moral value contained in each possible outcome where pressing the button is wrong is easy to calculate.

- (i) **Press and Wrong:** If one presses the button and it is wrong to press the button, then pressing the button produces less net happiness than the alternative; one's action has significant negative moral value.
- (ii) **No press and Wrong:** If one refrains from pressing the button when pressing the button would be wrong, then one's action produces more net happiness than the alternative; one's action has significant positive moral value.

Turn next to the two outcomes where pressing the button is not wrong. There are two ways for pressing the button not to be wrong: on one, it is because pressing the button produces the same amount of net happiness as not pressing (and hence both actions are permissible); on the other, it is because pressing the button produces more net happiness (and hence is obligatory). The value of one's action is as follows in each of these two sub-cases:

**(iii) Press and Not wrong:**

- (a) If it is not wrong to press the button because pressing and not pressing have the same moral value, then they produce the same amount of net happiness.
- (b) If it is not wrong to press the button because pressing produces more net happiness than not pressing, then pressing has significant positive value.

**(iv) No press and Not wrong:**

- (a) If it is not wrong to press the button because pressing and not pressing have the same moral value, then they produce the same amount of net happiness.
- (b) If it is not wrong to press the button because pressing produces more net happiness than not pressing, then refraining from pressing has significant negative value.

Overall, if pressing is not wrong and one presses, then (since one cannot know whether it is not wrong because one is in a case of type iii.a or iii.b), pressing has a small positive value. And, if pressing is not wrong and one does not presses, then (since one cannot know whether it is not wrong because one is in a case of type iv.a or iv.b), not pressing has a small negative value.

All of this can be summarized in the following table for each borderline case  $b_i$  (as before,  $\mathbf{b}_i$  is the proposition that pressing the button in  $b_i$  is wrong;  $\neg\mathbf{b}_i$  is its negation):

	$\mathbf{b}_i$	$\neg\mathbf{b}_i$
press	large -	small +
no press	large +	small -

With this decision table in hand, we can give some schematic answers to the question of what one rationally ought to do in a case of ethical vagueness. One ought to maximize expected moral value, where the expected moral value of pressing the button and not pressing the button in a borderline case is calculated by multiplying the values of the possible outcomes of an action by their likelihood.

Since one rationally ought to press the button just in case the expected moral value of pressing is greater than the expected moral value of not pressing, this amounts to the following. Where  $Val_p$  is the moral value of pressing if  $p$  is true, one ought to press in a borderline case  $b_i$  just in case the following holds:

$$Val_{b_i}(\text{press}) \times Cr(\mathbf{b}_i) + Val_{\neg b_i}(\text{press}) \times Cr(\neg \mathbf{b}_i) > Val_{b_i}(\text{no press}) \times Cr(\mathbf{b}_i) + Val_{\neg b_i}(\text{no press}) \times Cr(\neg \mathbf{b}_i)$$

Even without being any more specific than ‘large positive value’, ‘small negative value’, etc. in specifying the values contained in the various possible outcomes, we can note an interesting feature of the decision problem.

Begin with the “middle” borderline case where the likelihood that pressing the button is wrong is .5. Why one ought to refrain from pressing the button is straightforward: in these cases, if one presses the button, there is an equal chance that one produces a large negative outcome or a small positive outcome. And if one doesn’t press the button there is an equal chance that one produces a large positive outcome or a small negative outcome. Thus refraining from pressing has higher expected moral value.

Given INCREASING CREDENCES and UNIFORMITY, there will be some cases  $b_i$  past the “middle” borderline case where the likelihood on what one knows that  $b_i$  is true is less than .5. In these cases, because of the asymmetry in the moral values under the possible outcomes of pressing being wrong in  $b_i$ , and pressing being not wrong in  $b_i$ , *one rationally ought to refrain from pressing the button* in these cases. This is so even though one thinks it more likely than not that pressing the button is not wrong.

One needn’t be rationally required to refrain from pressing the button in every borderline case. In some cases one’s confidence that pressing the button is not wrong is so great it will make  $Val_{b_i}(\text{press}) \times Cr(\mathbf{b}_i) + Val_{\neg b_i}(\text{press}) \times Cr(\neg \mathbf{b}_i)$  exceed  $Val_{b_i}(\text{no press}) \times Cr(\mathbf{b}_i) + Val_{\neg b_i}(\text{no press}) \times Cr(\neg \mathbf{b}_i)$ . Where exactly these cases lie will depend on how big the difference in the moral values in the above table are. Since we haven’t supplied the needed detail, we cannot take a stand on a more concrete recommendation here. But the structural feature still holds: *some* cases where one’s credence that pressing the button is not wrong is greater than .5 are still cases where one rationally ought to refrain from pressing.

There are a number of assumptions in the foregoing: Utilitarianism about moral obligation and Epistemicism about vagueness are two. I discuss what happens when we discard these assumptions in the next sections. But before moving to these issues we can also note that, with these assumptions in place the same kind of reasoning is required when there are no cases of type iii.a and iv.a, where one’s options being of neutral value is possible.

Here is one way (given Epistemicism and Utilitarianism) for it to be vague what one ethically ought to do, but for it to be knowable that one’s action is not of neutral moral value. Take the modified sorites series in §1 where in each case, one can kill 1,000,000 creatures in order to save an adult human, but in each borderline

case, it is vague whether the 1,000,000 creatures in each can experience happiness. Given Epistemicism, each borderline case is a case where one can't know whether the 1,000,000 creatures can experience happiness, and hence can't know whether pressing the button and killing the 1,000,000 creatures is wrong. Nonetheless, in each borderline case it will either be true that pressing the button produces vastly more happiness, or produces vastly less happiness, than the alternative. So there are no borderline cases where it is true (though unknowable) that pressing the button produces the *same* amount of happiness as not pressing.

Nonetheless our decision table will have the same structure as before, and hence the "somewhat conservative" recommendation of the rational 'ought' will still apply. To see this, we can again divide a borderline case  $b_i$  into four possible outcomes. In the outcome where one presses the button and pressing is wrong, one saves the life of a human (say, a gain of +1,000 units of happiness) but kills 1,000,000 happiness-experiencing creatures (say, a loss of -1,000,000 units of happiness). When one presses when pressing isn't wrong, one gains the life of the human (+1,000) and loses nothing. When one refrains from pressing when pressing isn't wrong, one loses the life of the human (-1,000) and gains nothing. And when one refrains from pressing when pressing is wrong, one loses the human (-1,000) but saves the 1,000,000 creatures (+1,000,000).

The decision table then looks as follows:

	$b_i$	$\neg b_i$
press	-999,000	1,000
no press	999,000	-1,000

Thus the decision table has the same structure as before, and the somewhat conservative recommendation of the rational 'ought' follows for the same reason.<sup>12</sup>

#### 4 Interlude: complications and connections

The foregoing section argued for a specific result for the deliverances of the rational 'ought' in cases of ethical vagueness that fit the pattern of a sorites series for 'wrong'. We assumed Epistemicism about vagueness and Utilitarianism about moral value in deriving these results. Some of these assumptions are innocuous, and can be discarded. Others are essential to the result. I briefly discuss some of these surrounding issues in this section, before turning to a class of theories of vagueness which determinately deliver different results for the cases at hand in §5.

##### 4.1 *Rankings of value*

For simplicity the preceding discussion equated the moral value of an action with the net amount of happiness it produces. Obviously not every moral theory will

<sup>12</sup>Thanks to Elizabeth Barnes for helpful discussion on this issue.

agree with this simple Utilitarianism. We might then worry that a ranking on outcomes which we helped ourselves to won't be straightforwardly available for the rational 'ought' outside of this context.

But here there is reason to be optimistic that the formal needs of the expected moral value maximization approach can be met outside the context of simple Utilitarianism. There is however some precedent in the ethics literature in support of the claim that any ethical theory meeting minimal constraints can be represented in the same form as the Utilitarian theory. This has been done in the work of "consequentializers" such as Dreier (1993), Louise (2004), and Portmore (2007), who derive something like the "moral value" contained in each outcome for standard deontological moral theories. More precisely, a non-Utilitarian moral theory can be associated with a ranking on actions. This ranking can then be used to determine an expected goodness of an action, once the probabilities are filled in.

Of course there are still complications here—if the consequentialized ranking is merely ordinal in structure, then the decision-theoretic framework will have to be revised. And the ranking may need to be relativized to parameters (times, worlds, etc.) that make it unsuited for the role of being what rationality requires expected maximization of.<sup>13</sup> But in principle the consequentializing project shows that the results here are not limited to simple Utilitarianism.

#### 4.2 *Non-Epistemicist approaches to vagueness*

So far we have only answered the practical question of what to do in a case of ethical vagueness under an assumption of Epistemicism about vagueness. This assumption is of course not uncontroversial and many will no doubt wish to reject it. There is no guarantee that the rational 'ought' will deliver the same verdicts if Epistemicism is false.

Here is one way in which relaxing this assumption can produce alternative recommendations from the rational 'ought'. Some views of vagueness do not answer key questions for the expected value maximization approach. These are questions about the likelihood that an action in a borderline case is wrong (the *credal* question), and the possible moral values of actions in border line cases (the *value* question). Epistemicism, I have argued, provides specific answers to these questions (the credal question, for instance, is partially settled by INCREASING CREDENCES and UNIFORMITY). But—to take one example—on linguistic approaches to vagueness inspired by Fine (1975) or Lewis (1982), where vagueness is semantic underdetermination or ambiguity, it is not clear what one's credences toward a borderline case should be. Since the credal question is naturally taken to be a question of what credence to have toward the proposition that pressing the button is wrong in a borderline case, and ex hypothesi it is underdetermined (or ambiguous) which proposition 'pressing the button is wrong' expresses, the credal

---

<sup>13</sup>Brown (2011)

question does not have an (obvious) answer on these views.<sup>14</sup>

A similar point goes for views which require that it be vague whether you believe that pressing that button is wrong (Dorr 2003), or views which explicitly require some non-credal attitude toward vague cases (Schiffer 2000). These views won't provide determinate answers to the question of what credence one ought to have toward vague borderline cases, and so will leave possible answers to the credal question underdetermined. If so, these views also will not settle which action in a borderline case maximizes expected moral value.

## 5 What to do when it's vague: non-Epistemicist views

One way in which a view of vagueness might disagree with the Epistemicist over of what to do in a borderline case is by not providing determinate answers to the credal or value questions. The end of the previous section sketch some views in this category. Here is another way: by disagreeing over which answer to the credal question is correct. This involves holding that something other than the Epistemicist's credences are required toward borderline cases. A third way is to disagree on the value question: this involves holding that the possible values of one's action in a borderline case are not as the Epistemicist says they are.<sup>15</sup>

I will discuss one example of each kind of view in this section. The discussion is by no means complete. But each example provides a concrete case for the conclusion that the rational 'ought' in cases of ethical vagueness does not operate independently of the nature of vagueness.

### 5.1 *Alternative credences: the non-classical case*

Hartry Field (2000, 2003) characterizes vagueness in terms of the appropriate credal states toward a borderline case. The central feature of the view is that credences in borderline cases ought to be non-classical: if  $b_i$  is a borderline case of wrongness, and  $\mathbf{b}_i$  is the proposition that pressing the button is wrong, then  $Cr(\mathbf{b}_i)$  and  $Cr(\neg\mathbf{b}_i)$  do not sum to 1. The difference between  $Cr(\mathbf{b}_i) + Cr(\neg\mathbf{b}_i)$  and 1 is greater to the extent that it is certain that  $b_i$  is a borderline case. That is, if one is highly confident that pressing the button in  $b_i$  is not determinately wrong, and highly confident that it is not determinately not wrong, then one will have a very low credence that pressing the button is wrong *and* one will have a very low credence that pressing the button is not wrong. Less indeterminate cases do not involve as substantial a departure from classical probability: when  $b_i$  is a "fringe" borderline case,  $Cr(\mathbf{b}_i) + Cr(\neg\mathbf{b}_i)$  will be closer to 1. And insofar as  $b_i$  is a clear borderline case,  $Cr(\mathbf{b}_i) + Cr(\neg\mathbf{b}_i) = 0$ .<sup>16</sup>

<sup>14</sup>Cf. Williams (2014: 393) on the "classical" version of supervenientism.

<sup>15</sup>Not every view needs to disagree with the Epistemicist on these questions; for instance metaphysical views about vagueness found in Barnes (2010), Cameron (2010), and Barnes and Williams (2011) might deliver the same answers.

<sup>16</sup>Field (2003: 466); Field (2000: 17) presents a slightly different version of the non-classical view.

To compare the recommendations of this view with the Epistemicist's, I will assume the following. First, some of the "borderline" cases on the Epistemicist view are clearly borderline on the non-classical view. (Credences for the non-classicist sum to 0 in these cases.) But second, other cases that are "borderline" for the Epistemicist are possibly borderline for the non-classicist, and so receive some non-0 non-classical credence.

Given these assumptions, UNIFORMITY is false on the non-classical view. Let  $b_{j-1}$  and  $b_j$  be a pair of borderline cases that are very close to a case where pressing the button is determinately not wrong, while  $b_{k-1}$  and  $b_k$  are very close to the "middle" of the series of borderline cases; they are almost certainly borderline. On the non-classical view,  $Cr(\mathbf{b}_{j-1}) + Cr(\neg\mathbf{b}_{j-1})$  will be very close to 1 (and similarly for  $b_j$ ), while  $Cr(\mathbf{b}_{k-1}) + Cr(\neg\mathbf{b}_{k-1})$  will not be anywhere near 1 (and similarly for  $b_k$ ). And so  $Cr(\neg\mathbf{b}_{j-1} \wedge \mathbf{b}_j)$  will be close but not equal to 1, while  $Cr(\neg\mathbf{b}_{k-1} \wedge \mathbf{b}_k)$  will be nowhere close to 1.

The difference in this approach to the credal question shows up in the rational 'ought' as follows. Take a simplified case where there are 100 (possibly or clearly) borderline cases; one's credences at the edges of the series must sum to close to 1, but not for middle cases. A partial representation of this might look something like the following:

	$Cr(\mathbf{b}_i)$	$Cr(\neg\mathbf{b}_i)$	$Cr(\mathbf{b}_i) + Cr(\neg\mathbf{b}_i)$
$b_1$	0	.9	.9
$b_{50}$	0	0	0
$b_{100}$	.9	0	.9

The middle case is a "clear" borderline case, so neither the claim that pressing the button is wrong nor its negation receive any positive credence in this case. Cases on the edge are possibly borderline, so  $Cr(\mathbf{b}_i) + Cr(\neg\mathbf{b}_i)$  will be non-0 (and, as the case is near a clearly non-borderline case, this sum will approach 1). One of these credences should be 0. Since one has credence 0 in case  $b_{50}$  in both the claim that it is wrong to press the button and its negation, one of these credences should stay at 0 as one moves away from the middle case. One shouldn't grow more confident that pressing the button is wrong as one moves toward  $b_1$  and further away from cases where pressing is determinately wrong. And one shouldn't grow more confident that pressing the button is not wrong as one moves toward  $b_{100}$  and further away from cases where pressing is determinately not wrong.

The non-classicist who accepts this probability distribution will disagree with the Epistemicist over what the rational 'ought' recommends in the cases she construes as definitely borderline, as well as some of the cases that are possibly borderline on her view.

Definitely borderline cases are cases such as  $b_{50}$ : since one has credence 0 in both possible outcomes, the expected moral value of either pressing or not pressing will both be 0. Either action is rationally permissible.



There are two kinds of a possibly borderline case for the non-classicist. In some of the possibly borderline cases, the claim that pressing the button is wrong gets positive credence, while its negation gets none. Not pressing the button will have higher expected moral value and one rationally ought to refrain from pressing. And in other cases the claim that pressing the button is not wrong gets positive credence, and one rationally ought to press the button.

So the non-classicist's recommendations are *symmetrical*, unlike the Epistemicist's: pressing and not pressing are recommended in the same proportion of borderline cases. And the non-classical recommendation is more *permissive* since there is a range of cases (the clearly borderline cases) where pressing and not pressing are ranked equally by the rational 'ought'.

## 5.2 *Alternative values: the degreed approach*

Other views of vagueness will answer the value question differently than the Epistemicist, by taking a different view on *how bad* it is to press the button in borderline cases. One way to take a stand on this question is to hold that, in cases of ethical vagueness, there are actions are wrong to a less-than-full degree. These are *degreed* approaches to vagueness.<sup>17</sup>

These views hold that borderline cases are characterized in the first instance by the fact that, in these cases, pressing the button is wrong to some degree that is less than 1, and greater than 0. (Here terminology differs between authors: one might say the *degree of truth* is between 0 or 1, or the *degree of determinacy* is between 0 or 1; or one might say that the *degree of wrongness* is between 0 and 1. I will not focus on the differences between these locutions in what follows.) Truth to degree 1 or 0 is a feature of determinate cases; borderline cases only qualify for truth of intermediate degrees.

Degrees of truth in borderline cases are structured so that pressing the button in those cases closer to cases where pressing the button is determinately not wrong is not wrong to a high degree. And pressing the button in borderline cases close to cases where pressing the button is determinately wrong is not wrong to a low degree. The degree of wrongness for cases in between will increase uniformly as cases approach a non-borderline case where pressing is determinately wrong.

It is natural to assume that degrees of truth are connected to the moral value of pressing the button in the following way:

**DEGREE-VALUE CONNECTION** The moral value of an action in a borderline case of wrongness  $b_i$  is a function of the degree of wrongness of that act in  $b_i$ .

Some degrees theorists endorse theses which are naturally related to DEGREE-VALUE CONNECTION; for instance Williams (2014) endorses the claim that if you are concerned only with your own welfare, then you should care about the welfare

<sup>17</sup>Examples include Edgington (1997), Smith (2008), and Williams (2014).

of a future individual to the degree that it is determinate that the individual in question is identical to you.<sup>18</sup>

Since moral value is the other dimension of an expected value calculation, the degrees theorist incurs concrete commitments for the deliverances of the rational ‘ought’ if she endorses DEGREE-VALUE CONNECTION. How bad it is morally to press the button in the borderline case  $b_1$  is settled by the degree to which pressing the button in  $b_1$  is wrong.

In a simple case where there are borderline cases  $b_1 - b_{100}$ , we can illustrate this by letting the value pressing the button in borderline cases be a simple function of the degree of wrongness in those cases. If  $d(\phi)$  is the degree of wrongness of an action,  $\bar{d}(\phi)$  is the degree to which  $\phi$ -ing is not wrong, the value of the action is  $Val(\phi)$ , then  $Val(\phi) = 2(\bar{d}(\phi) - .5)$ . This yields the following partial table of the moral values (where possible values range from 1 to -1):

	$\bar{d}(\text{press})$	$Val(\text{press})$	$Val(\text{no press})$
$b_1$	.99	.98	-.98
$b_{49}$	.51	.02	-.02
$b_{50}$	.5	0	0
$b_{51}$	.49	-.02	.02
$b_{100}$	.01	-.98	.98

This is a natural application of the degrees theory, but there are arguments in Edgington (1997) which suggest that she would reject DEGREE-VALUE CONNECTION.

In general, if vagueness involves degrees of truth, then there are fine-grained possible outcomes one can have preferences over. If there is vagueness in whether something is  $F$ , then one’s preferences might rank not only  $F$ s and non- $F$ s, but also things that are  $F$ -to-degree-.5,  $F$ -to-degree-.7, etc. Edgington point, which might appear to be inconsistent with DEGREE-VALUE CONNECTION, is that someone who prefers (determinate)  $F$ s to (determinate) non- $F$ s might nonetheless prefer non- $F$ s to things that are  $F$ -to-degree-.5. Thus she might rationally choose a determinate non- $F$  when choosing between a vague  $F$  and a determinate non- $F$ . And in general preferences between non-degreed outcomes does not force any structure on preferences over middling degreed outcomes. Edgington sums up: “Verities [i.e., degrees of truth] do have a role to play in a more refined account of decision, in giving, when relevant, a more fine-grained specification of the possible states of affairs over which our preferences and credences range.”<sup>19</sup>

This might naturally suggest that the above table of the moral value of actions is not a commitment of the degrees theorist (nor does it approximate one), and

<sup>18</sup>Williams (2014: 405), in the context of a discussion of indeterminate survival cases from Parfit (1971).

<sup>19</sup>Edgington (1997: 313). Thanks to an anonymous referee here.

more generally that DEGREE-VALUE CONNECTION has to be false. But this quick impression would be incorrect. The schematic case involving *F*s and non-*F*s relies on preferences to rank the value of a non-*F* above the value of something that is *F*-to-degree-5. This is, however, not at issue in decision-making in cases of ethical vagueness. What settles the moral value of an act is not one's preferences directly over outcomes. Since, if Edgington is right, it is only in cases where one has preferences that do not align with degrees-of-truth that the value of outcomes does not correspond to which degreed properties, Edgington's position is not inconsistent with DEGREE-VALUE CONNECTION.

So far we have addressed only the issue of the moral value of an action in an outcome for the degrees theorist. This is just one side of an *expected* moral value calculation. We also need an answer to the credence question.

The most natural perspective on the credal question holds that one needn't be uncertain about anything in vague cases. (I will address one complication to this natural picture below.) The vagueness is accounted for by less-than-full degrees of truth; in principle one needn't be uncertain at all as to what the facts about degrees of truth are in a borderline case are. On this approach, one is rationally required to *not be conservative at all* in borderline cases of wrongness. The rational 'ought' requires one to press the button in each of  $b_1 - b_{49}$ , allows one to either press or not press the button in  $b_{50}$ , and requires one not to press the button in  $b_{51} - b_{100}$ .

Note that, while uncertainty plays no role in the expected moral value calculation on this view, the rational 'ought' is not thereby equivalent to the ethical 'ought'. Vagueness induces less-than-full degrees of truth. Degrees of truth settle an answer to the value question: how bad is it to press the button in various possible outcomes? And the degrees-of-truth answer to this question which follows from DEGREE-VALUE CONNECTION yields a determinate verdict as to which action maximizes expected moral value. So the rational 'ought' will determinately apply in cases of ethical vagueness, and hence it will be true to degree 1 or 0 that one ought to press the button in each case. Since the degree of ethical wrongness is not 1 or 0 in these cases, what one is rationally required to do and what one is ethically required to do are not the same thing.

There is one complication to the above picture, which deserves to be explored further but I will only briefly mention here. Some degrees theorists hold that in addition to vagueness giving rise to degrees of truth, vagueness *also* requires some less-than-full belief. One should, on these views, proportion one's degree of belief that pressing the button is wrong to the degree to which pressing is wrong. That is, if pressing the button is wrong to degree  $n$  in case  $b_i$ , then one's credence that pressing the button in  $b_i$  is wrong should be  $n$ . We can call this thesis DEGREE-CREDCENCE CONNECTION.<sup>20</sup>

DEGREE-CREDCENCE CONNECTION and DEGREE-VALUE CONNECTION do not apply to the same entities. The latter thesis requires that values are proportioned to

<sup>20</sup>This position is endorsed in Smith (2008) and Williams (2014).

*degreed* propositions, which include the proposition that pressing the button in case  $b_i$  is wrong to degree  $n$ . And the former thesis constrains credences in propositions about wrongness simpliciter, not degreed propositions. There are several approaches one could take in giving an account of how DEGREE-CREDENCE CONNECTION and DEGREE-VALUE CONNECTION interact in settling what the rational ‘ought’ recommends.

One could hold that both theses have the same consequences for the rational ‘ought’. Since we know what DEGREE-VALUE CONNECTION recommends—it requires pressing the button in all of the cases where pressing is not wrong to a degree greater than .5—it follows that credences from DEGREE-CREDENCE CONNECTION must feature in an expected value calculation which is such that the following holds iff  $\bar{d}(\text{press}) > .5$ :

$$Val_{\mathbf{b}_i}(\text{press}) \times Cr(\mathbf{b}_i) + Val_{\neg\mathbf{b}_i}(\text{press}) \times Cr(\neg\mathbf{b}_i) > Val_{\mathbf{b}_i}(\text{no press}) \times Cr(\mathbf{b}_i) + Val_{\neg\mathbf{b}_i}(\text{no press}) \times Cr(\neg\mathbf{b}_i)$$

(Recall that  $Val_p$  gives the moral value of an action if  $p$  obtains, and  $\mathbf{b}_i$  is the proposition that pressing is wrong simpliciter in  $b_i$ .) This is what Williams (2014: §2.4) calls a “reconciliation project” in a related context.

Alternatively one could hold that credences that are estimates of degrees of truth arising from vagueness do not interact with the rational ‘ought’ in the way credences arising from ordinary uncertainty do. This resembles the position in Smith (2008: 241), where he says that betting with credences arising from vagueness is ipso facto irrational.

These are interesting issues, but I will not pursue them further here since they do not bear on the main thesis of this section. On either approach, the degrees theorist will accept a rational ‘ought’ that maximizes expected moral value, and in some cases recommends actions that are not those the Epistemicist recommends. Rational action in borderline cases is not an issue we can theorize about absent some reasonably concrete view about what vagueness is.

## 6 Concluding remarks

The forgoing sections explore some answers to the question, what does the rational ‘ought’ require in cases of ethical vagueness? Candidate answers divide into two types: ones that can deploy an expected moral value maximization approach to the rational ‘ought’, and those that cannot. And we have seen that not all of the answers of the second type agree about what one rationally ought to do in cases of ethical vagueness. In closing I will point to some issues that deserve to be explored further.

Some questions concern views discussed in §4.3 which do not provide determinate answers to the question of which action in a borderline case maximizes expected moral value.

First, such theories might live with the existence of indeterminate requirements in these cases. On such views, when it is vague whether pressing the button is ethically wrong, then it is thereby also vague whether it is rationally required to press the button. Whether this is a plausible route, when other candidate views of vagueness do offer determinate answers, needs defending.

Second, the expected moral value maximization framework we have used for settling the requirements of a rational 'ought' might be discarded. This is attractive, but not required. An approach to the rational 'ought' which does not depend on the expected moral value framework deserves to be explored here—especially for theories that do not wish to take the first option outlined above. Some approaches to vagueness have been criticized for their revisionary consequences for classical logic (Williamson 1994), and we might ask the analogous questions about the consequences of a theory of vagueness for the decision-theoretic approach to practical reasoning.

A final point about the §5 views that can deploy the expected moral value maximization approach. As different theories under this umbrella will entail divergent answers to the question of which actions are best, given what one knows, we can ask whether some of these answers are intuitively more plausible than others. Since the rational 'ought' will recommend different actions in these cases, there are in principle grounds for comparing the theories. Ethical vagueness can in principle be a useful test case for substantive theories of vagueness.

## References

- Elizabeth Barnes. Ontic Vagueness: A Guide for the Perplexed. *Noûs*, 44(4):601–627, 2010.
- Elizabeth Barnes and J. Robert G. Williams. A Theory of Metaphysical Indeterminacy. *Oxford Studies in Metaphysics*, vol. 6, pages 103–148, 2011.
- Jeremy Bentham. *An Introduction to the Principles of Morals and Legislation*. 1781. Citations to the 1996 edition published by Oxford University Press.
- Campbell Brown. A New and Improved Supervenience Argument for Ethical Descriptivism. In Russ Shafer-Landau, editor, *Oxford Studies in Metaethics*, vol. 6, pages 205–218. Oxford University Press, 2011.
- Ross Cameron. Vagueness and Naturalness. *Erkenntnis*, 72:281–293, 2010.
- Cian Dorr. Vagueness without Ignorance. *Philosophical Perspectives*, 17:83–113, 2003.
- Tom Dougherty. Vague Value. *Philosophy and Phenomenological Research*, pages 1–21, 2013.
- James Dreier. Structures of Normative Theories. *The Monist*, 76:22–40, 1993.

- Dorothy Edgington. Vagueness by Degrees. In Rosanna Keefe and Peter Smith, editors, *Vagueness: A Reader*. MIT Press, 1997.
- Hartry Field. Indeterminacy, Degree of Belief, and Excluded Middle. *Noûs*, 34(1): 1–30, 2000.
- Hartry Field. No Fact of the Matter. *Australasian Journal of Philosophy*, 81(4):457–480, 2003.
- Kit Fine. Vagueness, Truth and Logic. *Synthese*, 30(3-4):265–300, 1975.
- Patrick Greenough. Vagueness: A Minimal Theory. *Mind*, 112:235–281, 2003.
- John Hawthorne. *Knowledge and Lotteries*. Oxford University Press, 2004.
- Rosanna Keefe. *Theories of Vagueness*. Cambridge University Press, 2000.
- Angelika Kratzer. What “Must” and “Can” Must and Can Mean. *Linguistics and Philosophy*, 1:337–355, 1977.
- David Lewis. Logic for Equivocators. *Noûs*, 16:431–441, 1982.
- Jennie Louise. Relativity of Value and the Consequentialist Umbrella. *The Philosophical Quarterly*, 54(217):518–536, 2004.
- Derek Parfit. Personal Identity. *The Philosophical Review*, 80:3–27, 1971.
- Douglas W. Portmore. Consequentializing Moral Theories. *Pacific Philosophical Quarterly*, 88:39–73, 2007.
- Jacob Ross. Rejecting Ethical Deflationism. *Ethics*, 116:742–768, 2006.
- Thomas Scanlon. *What We Owe to Each Other*. Belknap Press of Harvard University Press, 1998.
- Stephen Schiffer. Vagueness and Partial Belief. *Philosophical Issues*, 10:220–257, 2000.
- Mark Schroeder. *Slaves of the Passions*. Oxford University Press, 2007.
- Andrew Sepielli. What to Do When You Don’t Know What to Do. In Russ Shafer-Landau, editor, *Oxford Studies in Metaethics*, vol. 4, pages 5–28. Oxford University Press, 2009.
- Russ Shafer-Landau. Vagueness, Borderline Cases, and Moral Realism. *American Philosophical Quarterly*, 32(1):83–96, 1995.
- Nicholas J. J. Smith. *Vagueness and Degrees of Truth*. Oxford University Press, 2008.
- Ryan Wasserman. Personal Identity, Indeterminacy, and Obligation. In Georg Gasser and Matthias Stefan, editors, *Personal Identity: Complex or Simple?* Oxford University Press, 2013.

J. Robert G. Williams. Decision Making under Indeterminacy. *Philosophers' Imprint*, 2013.

J. Robert G. Williams. Nonclassical Minds and Indeterminate Survival. *The Philosophical Review*, 123(4):379–429, 2014.

Timothy Williamson. *Vagueness*. Routledge, 1994.