Vagueness, Normativity, and Inescapable Questions

Rohan Sud[†]

The concept BALD displays all of the features paradigmatic of vagueness. For instance, it is *sorites-susceptible*. We can construct a series of scalps where the first member has no hair and each member has one more hair than the previous member. On the one hand, early members of the series are clearly bald and later members are clearly not bald. Yet we feel compelled to make the paradoxical soritical inference from the premise that *such-and-such member of the series is bald* to the conclusion that *the next member is also bald*. Pretheoretically, the inference appears unassailable.

Moreover, the concept allows for cases of a distinctive sort, ones that are located in the middling part of the sorites series. As a rough initial gloss: these are cases for which it seems "there is just no fact of the matter" whether the member is bald. I take it that such judgments are familiar. They are the ones we learn to express with semi-technical expressions like "it's borderline/vague/indefinite/indeterminate" whether the member is bald. And we learn to call these "borderline" cases.

Now consider instead our thinnest deontic concepts, for instance the moral OUGHT or — even thinner — the subjective, all-things-considered OUGHT. These ought-concepts display some of the symptoms emblematic of vagueness. For instance, they are plausibly sorites-susceptible. Imagine Regina has promised to attend a friend's party. On her way to the party, a stranger asks for assistance that would require her to miss the party. We can construct a series of cases beginning with cases where the stranger's need is trifling and ending with cases where the stranger's need is serious, but where adjacent cases differ to a vanishingly small degree. At the start of the sorites, she presumably should keep her promise; at the end of the sorites, she presumably should help the stranger. Yet a sorites inference is

[†]Penultimate draft. Forthcoming in *Philosophical Studies*. Thanks to Mercedes Corredor, Aleksander Domoslawski, Miguel Dos Santos, Matti Eklund, Daniel Hoek, Claire Kirwin, Gideon Rosen, Erica Shumener, Sharon Street, Zach Thornton, Kelly Trogdon, and JRG Williams for helpful conversations. Special thanks to Martín Abreu Zavaleta and Singa Behrens for extended comments. Thanks also members of Abreu Zavaleta's Vagueness seminar at Syracuse University, the Vagueness and Ethics Workshop at Uppsala, and an audience at the 2024 Chapel Hill Normativity Workshop.

extremely tempting: there is a strong pull to infer from the claim that *Regina ought to keep her promise in this case*, to *she ought to keep her promise in the next case*.

But consider the middling cases in this sorites series, where we would expect to find borderline cases. As many authors have confessed, the middling cases in these sorts of normative sorites *just feel different* from the corresponding cases in sorites series for more prosaic concepts.[1] That's to say, our judgments of such cases occupy a role in our broader mental lives that is very different from the role played by judgments of typical borderline cases. It's hard to articulate that difference, and different authors have offered different characterizations. As a first stab we can take our cue from Wasserman (2004, 400) and try saying that such cases are not accompanied by the full blown "no fact of the matter" phenomenology that we find for quintessentially vague concepts.

Consider a borderline bald man – one who is just "kinda sorta" bald. There doesn't seem to be a fact of the matter as to whether or not he is indeed bald. For instance, it would be absurd to ponder the matter. But now put yourself in Regina's shoes as she bumps into a stranger — one with a "kinda sorta" important need — asking her for help on her way to a friend's party. It would be normal for Regina to act as if there is some fact about what she should do, albeit one which is very hard for her to come to grasp. For instance, it would be reasonable for Regina to ponder whether she should help the stranger or not. And she might consider whether she acted as she should have long after she's decided one way or another. Unlike in the bald-series, there seems to be some hidden fact of the matter, one that we can reasonably be anxious to discover.

Unfortunately, all this talk of there seeming to be or not be "a fact of the matter" isn't particularly well-suited for careful theorizing. For now though, just treat it as a helpful label for a yet-to-be-characterized psychological difference. Indeed, let's go ahead and add a few more helpful labels. Let's say that these ought-sorites contain *quasi-borderline* cases, where a quasi-borderline case is one that lacks that nebulous "no fact of the matter" phenomenology, but has the other features of borderline cases (e.g. arises in the middle of a sorites). We'll reserve the standard term *borderline case* for cases that display the full suite of vagueness-related features, including that nebulous "no fact of the matter" phenomenology. And we'll use the

[1]For relevant discussion see: Chang (2002, §3), Wasserman (2004, 400), Dougherty (2013, 370), Schoenfield (2016, 265-6), Constantinescu (2012, 62), Williams (2016), Elson (2017, §6), Liao. (To be clear: some of these authors will disagree with my characterization of the phenomenon.)

neutral term *middling case* to refer to both quasi- and full-blown borderline cases.

With the help of this terminology, we can state the central puzzle of this paper as two questions.

(1) Why do sorites series built around descriptive concepts allow for full-blown borderline cases, and not merely quasi-borderline cases?

(2) Why do (at least some) sorites series built around the thinnest deontic concepts only allow for quasi-borderline cases instead of the full-blown borderline cases we would expect?

Now we have names for quasi-borderline cases and the unique phenomenology that accompany them. But what's in a name? Not much. My first task in this paper will be to propose a more rigorous characterization of the quasi-borderline / borderline distinction — a characterization that does (at least some) justice to that nebulous and pretheoretic "no fact of the matter" phenomenology.

Inspired by some observations made by Hartry Field (2000, 11-12) (2009, §11.2), Ruth Chang (2002, 684), and others, I will propose that we understand the quasi-borderline / borderline distinction in *zetetic* terms. Roughly: it's typically inappropriate to take a questioning stance (e.g. curiosity or wonder) towards middling cases in a descriptive sorites, but not in our puzzling ought sorites. That is, when it's borderline whether ϕ, the question whether ϕ is *closed*. That's not so for merely quasi-borderline cases.

With this proposed regimentation, our initial puzzle becomes a much more tractable puzzle about our *zetetic reasons*. Just as conceptual ethicists and engineers ask which concepts we should deploy and which we should replace, we can ask which *questions* we should deploy and which we should replace. Why do we typically have reason to close middling questions involving vague concepts, and what's special about thin deontic concepts that they resist such closure?

This, I claim, is a particularly fruitful way to think though our initial puzzle. It gives us a rare opportunity to theorize about vagueness and meta-ethics in tandem and holds out the potential to teach us lessons about each subject which we might miss if we theorize about them in isolation. I'll try to convince the reader that such lessons are indeed in the offing by giving my own take on the zetetic puzzle.

I'll start by addressing part (1) of our puzzle. After arguing that more popular theories of vagueness are ill-equipped to explain our target phenomenon, I'll offer my own explanation for why we typically close middling

questions. My approach builds on Eklund's (2002; 2019) *incoherentist* theory of vagueness, according to which vague concepts are governed by incoherent inference rules. Extending this claim, I'll argue that this incoherence gives us standing reason to refuse to ask certain question involving vague concepts. That's because refusing to ask such questions helps quarantine the incoherence inherent in such concepts.

Next, I'll turn to part (2) of our puzzle, regarding the normative sphere. There, I'll argue that the standing reason to refuse to consider middling questions is overridden. Given the role that normative concepts play in our practical reasoning, such questions are *inescapable* for practically deliberative agents like us. That's why ought-questions persist in the face of vagueness.

1 Refining the Puzzle

Compare our opening bald sorites-series and our ought sorites-series. We noticed that the series were similar in that they both exhibited sorites susceptibility. But we also noticed that our judgments of the middling cases in the two series had different characters. As a first (and very rough) approximation, I said that the judgments we direct at the middling cases in the bald-sorites came with a "no fact of the matter" feeling which is absent in the normative case. But I cautioned against reading any substantive metaphysics into this label; it's nothing more than a placeholder for an ill-defined psychological syndrome. Before we can theorize about this syndrome, we'll first need to characterize it more rigorously.

Let's do so by examining the distinctive roles that judgments of borderline cases play in our broader mental lives. Focus on the middling cases in the bald sorites, and the judgment we learn to express with locutions like "it's borderline whether he is bald".

Plausibly, such judgments exhibit a particular epistemic profile: when we decide that a case is borderline, we are taking it to be a case of *irremediable ignorance*. That is, we take the case to be one where we cannot know whether or not the member is bald despite knowing all of the baldness-making features of the case (e.g. the number and distribution of hairs).

But the phenomenology of such cases must go beyond our recognition of irremediable ignorance. As Field (2009, §11.2) points out, there are all sorts of things that we aren't able to come to know, but where recognizing that we are destined for ignorance doesn't produce the distinctive "no fact of the matter" phenomenology that we find in borderline cases. For instance, as he notes, we aren't able to learn facts about the inner workings of black

holes or the seventeenth significant digit of the Centigrade temperature at the currently hottest point in the interior of the sun. But, of course, these cases lack the "no fact of the matter" phenomenology.

The same thing applies to the quasi-borderline cases in our normative sorites. Plausibly, these are also cases of irremediable ignorance. In the middling segment of the series, it becomes unclear whether or not Regina ought to break her promise, even though we seem to know all of the rightness- and wrongness-making features of the case. But, as we've been saying, the phenomenology of these cases simply doesn't match that of borderline cases — it lacks that "no fact of the matter" feel, patterning instead with e.g. Field's temperature of the sun cases. So, there must be more to the "no fact of the matter" phenomenology accompanying judgments of borderline cases than irremediable ignorance.

Perhaps the phenomenology has more to do with the bouletic profile of borderline cases. Here's what I mean. Andrew Bacon (2018, §10.2) has claimed that vague propositions rationally require a certain sort of bouletic indifference. To a first approximation, he claims:

Bacon's Principle (Rough) It's irrational to intrinsically desire vague propositions.

A bit more precisely: he claims that, conditional on the strongest precise propositions, our preference relation should be indifferent to any further vague propositions. For example, once we've counted the number of hairs on Harry's head and deemed him a borderline case, our desires should be indifferent as to whether or not he is *also* bald. That is, we shouldn't prefer the case of *his having exactly n hair and being bald* to the case of *his having exactly n hair and being not bald*.

Might this help explain why we feel that there isn't any fact of the matter as to whether Harry is bald? Perhaps when our preferences are wholly indifferent as to whether p, we feel as if there is no fact of the matter whether p. Then (assuming that we're rational) Bacon's Principle would explain why (once we've counted the hairs on his head) there seems to be no fact of the matter as to whether Harry is bald.

But just as our inability to know some proposition doesn't generate the "no fact of the matter" phenomenology, nor does our preferential indifference to some proposition. Consider various highly theoretical propositions, such as spacetime substantivalism or Goldbach's conjecture. Speaking for myself, my preference relation is entirely indifferent as to the reality of spacetime or the structure of the natural numbers. But these cases obviously lack the "no fact of the matter" phenomenology. In such cases, unlike

the baldness cases, there seems to some hidden "fact of the matter," one that I might, for instance, be curious about.

Upshot: the epistemic and bouletic profile of borderline cases are red herrings. Notice, however, that borderline cases also play a distinctive *zetetic* role. Suppose, for each person x in the bald-sorites, you ask yourself: is x bald? At the start, the answer is clearly "yes". At the end, the answer is clearly "no". What about for borderline cases? Suppose Harry is a borderline case of baldness. Now consider the question:

(Q1) Is Harry bald?

As we've already said, we don't know whether or not Harry is bald, even when we seem to know all of the baldness-making features of the case. But the phenomenon runs deeper. As Field (2000, 11-12; 2009, §11.2) notes, there is something odd about wondering whether he is bald. It's similarly odd to take or pursue other "questioning" attitudes or activities: it's incoherent to accept that he is borderline bald but continue to be *curious* whether he is bald or to *speculate* as to whether he is bald. In a word, your judgment that:

(AB) It's borderline whether Harry is bald.

closes the question you started with. Let's say that someone who raises, is curious about, speculates over, or wonders about a question is taking a *questioning stance* towards that question. To say that (AB) closes (Q1) is to say that accepting (AB) and taking a questioning stance towards (Q1) evinces conceptual confusion.

Now, the quintessential way we close questions is by coming to believe one of the answers to the question.[2] That is, the usual way we close the question of whether someone is F is by coming to believe something like:

(AY) Yes, they are F.

or

(AN) No, they are not F.

What we've seen is that soritical concepts allow for a third way to close questions — by judging a case to be borderline.

If you don't immediately find yourself agreeing that (AB) closes (Q1), two clarifications might help.

[2]See e.g. Friedman (2019, §3).

First: The incoherence claim I am making is widescoping. It is incoherent to *both* accept (AB) *and* take a questioning stance towards (Q1). If we suspend our judgment in (AB), we can *then* wonder / be curious / ask whether he is bald. Now, because it's often tricky to decide whether a particular case is indeed borderline, it's often perfectly reasonable to suspend our judgment in (AB) and then go on to pursue (Q1). That can make my claim harder to assess. One way to sidestep this issue is by considering the slightly different question:

(Q2) Who is the last bald member of the series?

(Q2) is akin to simultaneously asking, for each member of the series, whether that member is the last bald member. Given an appropriate series, it's very difficult to not admit that it's borderline who the last bald member is — even if we are less confident that such-and-such particular member is borderline bald. Here again, though, we clearly see the phenomenon of question-closure. Someone who takes a questioning stance towards (Q2) is evincing a sort of conceptual confusion, precisely because it is so obviously borderline who that last member is.

Second: I am not saying that we never exhibit this sort of confusion. Indeed, part of what makes sorites series so puzzling is that they can be used to get us to take up questions that we are otherwise inclined to keep closed. This is most obvious in the case of a *forced march sorites* in which an examiner pressures a subject to consider whether a particular member of the series is bald, and (immediately following a verdict) asks the subject whether the very next person in the series is bald. Importantly, the subject's only options for reply are "yes" or "no" (i.e. there is no option to respond with a shrug or with "don't ask me that!"). In such circumstances, a subject might genuinely consider and wonder about (Q2) or every instance of (Q1), especially at the moment they catch themselves swapping their responses from "yes" to "no". Similarly, philosophers of a certain bent (maybe those who have drunk too much of the epistemicist Kool-Aid!) might lead themselves down a forced march sorites, and they might thereby come to wonder about (Q1) and (Q2).

But surely such a questioning stance is abnormal — it is not our "default" state. We might be moved from our default state as the result of being prodded down a sorites, with pressure being applied (externally or internally) to consider questions that are otherwise closed. But that state is confused, just like when, in the grip of soritical reasoning, we issue confused and confusing verdicts like "Harry is both bald and not bald". So it's unsurprising that we return back to our default state of closure. Even if we temporarily

leave open the question of who the last bald man in the sorites is, the "no fact of the matter" phenomenology inevitably returns upon reflection, and we proceed to close it. So my claim is really more qualified: someone *whose default state includes* a questioning stance towards such questions is evincing a conceptual confusion.

What makes this observation so puzzling is that questioning stances are extremely permissive. This is especially true for the attitude of curiosity. As noted earlier, there is nothing incoherent about being curious whether p when you have no chance of coming to know whether p or when our desires are indifferent with respect to p. So why is it so confused to be curious whether Harry is bald?

The puzzle gets deeper when we consider the analogous cases in the normative sorites. Imagine we ask:

(Q3) Should Regina help the stranger rather than go to the party?

Imagine further that we learn some facts about Regina's case. For instance, we learn that she is in the middle part of a sorites. And we realize that her case suffers from irremediable ignorance: we've noted all of the relevant considerations of the case and are still unsure what she should do. In this case, I claim, the question remains open. We can continue to be curious or speculate about what Regina should do — even if we realize that we will never discover the answer. There is no requirement that our default state be one in which (Q3) is closed. The same is true with the question:

(Q4) When does it become permissible to break her promise and help the stranger?

There is nothing conceptually confused about someone who's default state is to be curious about this question. That's true even if they recognize that they will never come to know the answer to this question. And, unlike the baldness case, the curiosity that we might induce by considering a sorites series is stable. There is nothing confused about someone who doesn't "return" to a default state of closing the question.

Other authors haven't framed the phenomenological difference of normative sorites in quite such explicitly zetetic terms. But my framing does follow one thread in some authors' descriptions — descriptions that sometimes use zetetic language. For instance, Ruth Chang (2002, 682-8) has claimed that normative questions are saddled with a "persistent perplexity," even when they are accompanied with the other features of borderline

cases (see also Williams (2016, 417-18, 426-7), who emphasizes this perplexity). To illustrate, Chang asks us to consider two cases. In the first, Jack is required to sort people into two categories — bald or not bald. When he comes across a borderline case of baldness — Herbert — he must resolve the case. Here, she says, he would be entirely comfortable resolving the matter arbitrarily by, e.g., flipping a coin. As she puts it:

> The resolution of a borderline case lacks what we might call "resolutional remainder": given all the admissible ways in which the case might be resolved, **there is no further question** as to how resolution should proceed. (684, my emphasis)

In this case, she says, the "perplexity" doesn't persist. And, she adds:

> If we add another player, Jill, who happens to sort Herbert into the "not bald" pile, Jack and Jill have no real disagreement; their "disagreement" is simply a clash of arbitrary decisions in the face of indeterminate application. (684)

Chang doesn't think this is so for evaluative comparisons. She makes her case especially vivid with the predicate 'morally better than'. Faced with a middling case of this predicate (e.g. it is morally better for Regina to keep her promise or help the stranger?), one would not feel comfortable sorting based on a coin flip. In this case, she says, the perplexity over whether which action is morally better would persist. And two people who sort the case differently would be having a substantive disagreement.

This can all be cashed out in terms of question closure. When — in the course of deciding who is and is not bald — we come across a borderline case, it's inappropriate to consider whether they are bald. The question disappears, so to speak. So, without deciding whether or not they are bald, we simply sort the case arbitrarily. And there is no disagreement with someone who arbitrarily sorts them otherwise. We haven't come to a different verdict on the question of their status as bald — rather both parties have simply closed the question. Not so in the normative cases. The question of which is morally better is open for us. So, for instance, we can wonder whether we are sorting the case correctly. And we can disagree with someone who, also wondering about the question, decides one way or another.

Summing up: I propose to regiment the otherwise obscure talk of the "no facts of the matter" phenomenology in terms of question closure. A key difference between cases about the temperature at the interior of the sun and borderline cases of baldness is that in the former, but not the latter, we feel that there is "some hidden fact of the matter" that we can wonder / be curious / ask about (even if our questioning is bound to end in ignorance).

In the latter, taking a questioning stance (at least without special prodding) is conceptually confused. Our default state is to close such questions. That's why we aren't left with the appearance of some hidden fact of the matter. The same can't be said for middling cases in our normative sorites: we don't close the question as to what Regina should do. In any given case, we can continue to be curious (or even wonder) about what she should do — even when there is no hope of satisfying that curiosity. That's the sense in which there appears to be some hidden fact of the matter that we can be anxious to discover. And that's puzzling.

Our puzzle can now be stated more carefully, in terms of question-closure. Why does question-closure pattern with sorites susceptibility and irremediable ignorance for descriptive concepts, but not for thin deontic concepts?

2 Clarifications and Assumptions

Before pursuing our puzzle in earnest, I need to flag several assumptions and clarifications.

First. I've been following Bacon (2018), Schiffer (2003), and others in assuming that vagueness lies in our propositions. Here's what I mean by that. Some propositions are vague and others are precise. The distinction is supposed to be intuitive. Propositions of the form *x is bald* are vague. And microphysical propositions like *all electrons are negatively charged* are precise.[3] And importantly, we can hold propositional attitudes towards these vague propositions. On this picture: when I (definitely) believe that Harry is bald, there is a particular proposition that I am (definitely) in the belief-relation to — the vague proposition *that Harry is bald*.

Second. I am not claiming that normative concepts never make for full-blown borderline cases, cases with question-closure. Surely we can find question closure in the midst of *some* sorites series constructed with *some* normative terms. Question-closure is particularly easy to find with evaluative concepts, e.g. generous or good, with application conditions that plausibly depend on a vaguely determined threshold along some scalar dimensions of value. Still, it's enough to generate the puzzle to note that in some normative sorites series question-closure fails to pattern in the way we would expect.

[3] Although, see Bacon (2018, §12.2) for an argument that such propositions are, surprisingly, not precise.

Failures of question closure are easiest to find in cases involving thin deontic concepts. This is particular so for the "subjective," "deliberative," or (the term I will use) "decisional," all-things-considered OUGHT. By the decisional all-things-considered ought, I mean the sort of ought that Allan Gibbard (2003) points us to when he talks about deciding that something is "the thing to do" in a given choice situation.[4] It's the sort of ought that we come to a view on when we act deliberatively, even if we have limited knowledge about, e.g., the effects of our actions. In this respect, it differs from the objective ought: when acting with limited knowledge, we often remain uncertain of what we objectively ought to do. Our puzzle is most stark when it comes to cases of vagueness about what we decisionally all-things-considered ought to do. When, in the middle of the sorites, Regina continues to ask what she should do, it's easiest to interpret her as asking about this decisional all-things-considered ought. This being so, I will focus on middling cases of this concept, and (unless otherwise noted) the term OUGHT should be disambiguated as picking out this particular ought-concept.

Nevertheless, I take it that other ought-concepts (e.g. moral or prudential oughts; objective oughts) and other normative concepts (e.g. the concept of BEING MORALLY BETTER) have conceptual ties to this decisional all-things-considered OUGHT. For example, we might say that to judge that one morally ought to do x is (*very* roughly) to judge that one (decisionally, all-things-considered) ought to feel guilt upon not doing x (GIBBARD, 1990, Ch. 3). And we can say that to judge that one objectively ought to ϕ is (*very* roughly) to judge that, were one made epistemically ideal, one (decisionally, all-things-considered) ought to ϕ (GIBBARD, 2005). So, given the tie to other ought-concepts, I am hopeful that what I say about the decisional all-things-considered OUGHT can be extended to explain analogous phenomenon that arise for these other normative concepts.

Third. I am open to the idea that normative concepts aren't the *only* sorts of concepts that preclude question-closure. For instance, the first-personal concept ME and the concept CONSCIOUS arguably also resist question closure (i.e. we can always wonder whether a developing fetus is conscious or whether I survive a case of fission). But, in so far as these are genuine exceptions, they are the exceptions that prove the rule. My goal here is to explain that rule and outline a strategy for explaining one kind of exception — a strategy that we can hope will generalize to the other exceptions.[5]

[4] Although I follow Gibbard in positing this ought-concept, I will not be taking up his controversial non-cognitivism.

[5] In §5.2, I draw from Hofweber's (2023) notion of an *inescapable* concept to argue that

Fourth. As Weatherson (2009, §4.1.3) points out, whether someone is seduced by the soritical reasoning in a normative series — whether they find the normative sorites genuinely soritical — will depend on their background normative views.[6] For instance, compare two theorists. The first, Prem, adopts a flat-footed exceptionless deontological theory according to which one always ought to aim to keep one's promises regardless of the consequences of doing so. The second, Theresa, adopts a threshold deontological theory according to which one ought to aim to keep one's promises unless doing so would cause a significant amount of suffering. Prem will not find our original promise-vs-suffering series soritical. He will think that, throughout, Regina should keep her promise. Theresa, on the other hand, *will* find the series genuinely soritical.

I'm addressing the puzzle to those of us like Theresa who adopt normative views (even just partial ones) that make for soriticality. What's puzzling is that even those of us who find the above series paradoxical can be curious about what Regina should do in any particular case in that series.

Lastly. Question-closure might not be the *only* puzzling difference between the middling cases in normative and descriptive sorites. For instance, recall Bacon's Principle, according to which we should not intrinsically desire the vague. As an exception to this principle, several authors have argued that we don't see rational indifference in middling normative cases.[7] Imagine that Regina is in a middling case, and chooses to honor her promise to her friend instead of helping her neighbor. According to these authors, she might reasonably desire that she acted permissibly. That raises another puzzle:

Puzzle of Bouletic Exceptionalism What explains Bacon's Principle and the apparent exception of normative propositions? Why is it typically irrational to intrinsically desire vague propositions? And, given this, why is it apparently rational to intrinsically desire vague normative propositions?

---

ought-questions (as opposed to baldness-questions) are essential for a certain sort of inquiry. Hofweber's notion bears some resemblances to Chalmers' (2011, §8) notion of a *bedrock* concept. And, interestingly, Chalmers suggests that the concepts ought, consciousness and me might all be cases of bedrock concepts.

[6] Weatherson notes that competent users of normative concepts can have radically different "parameters of application" for those concepts. He makes this observation in the course of arguing against Eklund's proposed "definition" of 'vagueness'. I discuss Weatherson's observation a bit further on p. 21.

[7] Jack Spencer, for instance, explores some such cases involving moral oughts. See his (2022, §7-8).

The Puzzle of Bouletic Exceptionalism is, I think, a real puzzle. And I will have something to say about it later. But the focus on this paper will be the zetetic puzzle. As explained above, without explaining question-closure, we simply aren't doing justice to the felt sense in which there is a "hidden fact of the matter" in normative, but not descriptive cases.

3 Dead Ends

Recall the first half of our puzzle: why do descriptive concepts allow for full-blown borderline cases and not merely quasi-borderline cases? Why do we see question-closure at all?

Here's an idea. Stretching back to at least Frege, there is a tradition of thinking that vagueness is a sort of defect in our representations of an otherwise precise world. Assuming vagueness lies in propositions, this would mean that vague propositions are defective. But a question is just a way of considering various propositions. So, perhaps, vague questions are similarly defective.

Different proponents of the "vagueness-as-defect" tradition give different glosses of the relevant notions of defectiveness. For instance:

- An epistemicist understands the defectiveness in terms of unknowability.
- A truth-gap theorist (e.g. supervaluationists) understands the defectiveness in terms of failures of bivalence.
- A non-classical logician (e.g. intuitionists) understands defectiveness in terms of failures of classical laws like the law of excluded middle (LEM).

Unfortunately, none of these glosses are promising for our present puzzle. Consider the view of an epistemicist. As we noted above, we often ask questions we have no hope of coming to know the answer to. The epistemicist would need some story about how the special source of ignorance that supposedly accounts for vagueness requires question-closure. I see no way to tell such a story.

Consider instead the truth-gap theorist. Arguably, when we take a questioning stance towards whether p, we are assuming that p is either true or false, and we're interested in which.[8] That would help explain why, for

[8]Thanks to Gideon Rosen for helpful discussion here.

instance, we don't take questioning stances towards propositions we realize have empty concepts or false presuppositions. (I don't wonder whether Mary stopped smoking once I decide that she never started smoking or that my concept MARY is empty.) According to some theorists, when it's borderline whether Harry is bald, it's neither true that Harry is bald, nor false that Harry is bald. That would give us a tidy explanation of question-closure.

An analogous explanation is available for for those that reject classical logic. On some theories, LEM fails in borderline cases. But, arguably, when we take a questioning stance towards p, we're presupposing $p \vee \neg p$. Again, this would give us a tidy explanation for question-closure.

Unfortunately, I doubt failures of bivalence or LEM hold the key to solving our puzzle. Of course, there are generic reasons — ones having nothing to do with our current puzzle — to retain bivalence and LEM in the face of vagueness. But there is also a more specific reason in the current context: it seems hopeless to use such failures as a basis for explaining the difference between the descriptive and normative cases. Think about how such an explanation would have to go. We would need to argue that, while bivalence and LEM fail to hold in descriptive domains, they do hold in the normative domains. But, if anything, the case for bivalence and classical logic is weaker in the normative domain! In other words, the proposed explanation of the first half of our puzzle would leave us with no room to explain the second half of our puzzle. While failures of bivalence and non-classical logic give us tidy explanations for question-closure in typical cases, those explanations are *too* tidy, making it difficult to explain the exceptions in normative cases.

Given this difficulty, we may be tempted to adopt (what I'll call) a *disjunctive strategy*. Proponents of this approach claim that the source generating descriptive sorites susceptibility and irremediable ignorance are largely unrelated from whatever generates the normative analogues. Instead, they offer piecemeal explanations of the data, typically with radical metaphysical assumptions.

I place Ruth Chang (2002) in the disjunctivist camp. Focusing on comparisons of value, she takes the persisting perplexity associated with (what we're calling) quasi-borderline value comparisons to be evidence that the phenomenon is quite distinct from vagueness. According to orthodoxy, when we are comparing two options with respect to F, one option can be more F than the other, less F than the other, or the two can be equally F. Chang argues that when it comes to comparative value judgments (e.g. *morally better than*) there is a fourth way they can be compared: the two

options can be *on a par*. Importantly, for two options to be on a par is not to be understood as the options being incomparable. Nor is it to be understood as it's being vague which comparative relations holds. Instead, the parity relation is meant to be a substantial posit in the metaphysics of value, a *sui generis* way two things' value might compare. Chang takes the persisting perplexity of such cases as evidence that whatever is giving rise to borderline cases for prosaic non-value-laden comparisons and predicates cannot be what's giving rise to quasi-borderline cases. This is supposed to clear the way for her *sui generis* relation to step in: such quasi-borderline cases are instances of parity.

I also place another theorist, Ryan Wasserman (2004), in the disjunctivist camp. Wasserman adopts a background theory of reference and vagueness according to which (i) some properties are more "natural" or fundamental — they do better at "carving nature at its joints" in David Lewis's (1983) sense; and (ii) vagueness for notions like baldness is explained in part by the unnaturalness of the candidate referents for our term 'bald'. With this background theory, he acknowledges that intuitions for the middling cases in moral sorites are different from analogous descriptive cases: whereas middling descriptive cases are accompanied by a "no fact of the matter" intuition, the same is not true in the moral case. Wasserman then claims that, in so far as want to vindicate this intuition, the only way we can do so is by treating evaluative properties — but not properties in the vicinity of *baldness* — as natural or fundamental. Our "fact of the matter" intuition in the normative case would therefore be tracking the existence of a hidden joint in moral reality. So, it would seem to follow, the explanation for the vagueness-related features displayed by the notion of baldness has virtually nothing to do with the apparently similar features displayed by evaluative comparisons — in the latter case, but not the former, reference is fully determinate and there is simply a hidden naturalness boundary.

I find the disjunctive strategy to be wanting. Start by considering Chang's implementation. The radical metaphysics of value underlying that proposal can be difficult to swallow. Speaking for myself, I lose my grip on what it means to claim that two things are comparable with respect to F while also claiming that neither is more F than the other nor are they equally F. What's more, there are powerful arguments in the literature demonstrating that a parity relation violates various extremely plausible principles in the "logic" of comparisons.[9]

Consider Wasserman's view. Both his theory of reference and the sug-

[9]See, for instance, Dorr et al. (2023).

gestion that normative properties are natural or fundamental are deeply controversial. And, more importantly, his approach raises just as many questions as it answers. In particular, we are left wondering why normative terms display the other vagueness-related features. Given that there is a hidden naturalness boundary, why are we attracted to the sorites inference in the first place? And why is the naturalness boundary hidden at all? In so far as there is a workable epistemology for naturalness facts, we need some extra story why our "naturalness detectors" malfunction in the normative case.

Upshot: we should prefer a more unified and less metaphysically-controversial strategy for explaining question-closure and its lack.

***

Here's our dilemma. On the one hand, we are seeking a unified explanation of the vagueness-related features that we see across the descriptive and normative spheres, features like sorites-susceptibility and irremediable ignorance. But on the other hand, we want our explanation to leave room to explain the difference we find across the spheres – in particular the lack of question-closure for our concept OUGHT. That's a fine line to toe. But it can be done.

The trick, I claim, is to focus on our reasons for pursuing certain questions. If the truth-gappers or non-classicalists were right, we would be unable to take questioning stances towards borderline questions (just as we are unable to wonder about questions we recognize contain e.g. empty name-concepts). But that makes it hard to see why we can do so in the normative case. On my view, it's not that we are *unable* to take questioning stances towards vague questions. Rather, it's just highly irrational for us to do so. The issue is one of our zetetic reasons. Vagueness creates standing reason for us to drop descriptive questions. And that reason, I'll claim, is overridden when it comes to ought-questions. In what follows, I will try to vindicate this approach.

My vindication takes two steps. Step one: I'll sketch a theory according to which we have standing reason to drop descriptive questions in the face of vagueness. That sketch will be preliminary and high-level, but it will hopefully given the reader a proof of concept for the foregoing zetetic-reason strategy. Step two: I'll argue that those reasons can be overridden in the case of ought-questions. Together, these two steps give us a working example of my proposed zetetic-reasons strategy.

Importantly, these two steps need not be accepted as a package-deal. For instance, a reader need not accept my theory of our reasons to close

middling baldness-questions in order to accept what I go on to say about our reasons to keep ought-questions open. The two steps have independent interest.

4 Step One: Incoherentism and Question Closure

My preferred explanation for question-closure builds on the incoherentist approach to vagueness. Pioneered by Matti Eklund (2002; 2005; 2019), with inspiration from Dummett (1975), incoherentists claim that:

Incoherentism Users of vague concepts are required, as a matter of competence with that concept, to have a default disposition to apply inconsistent inference rules (where inference rules are *inconsistent* when one can derive a contradiction from repeated applications of those rules).

On this view, the sorites paradox built around a given concept is an illustration of the inconsistency of the rules that any competent user of that concept must use to regulate their inferences. As a toy example, an incoherentist might take BALD to be associated with the following inference rules:

R1. $\approx$ HAVING 0 HAIRS MAKES YOU BALD

R2. $\approx$ HAVING 50K HAIRS MAKES YOU *NOT* BALD

R3. All instances of: HAVING n HAIRS MAKES YOU BALD $\approx$ HAVING $n+1$ HAIRS MAKES YOU BALD

It's easy to see that, from repeated applications of these rules, we'll end up concluding contradictions like *having 50k hairs makes you bald and not bald*.

Importantly, this inconsistency is supposed to be "built into" the vague concept itself in the sense that the default state of anyone competent with the concept BALD will include a disposition to apply the inconsistent inference rules. The inconsistency is analytic, so to speak. The same is not true with other cases of inconsistent belief. For instance, suppose someone believes that *Marco is in the room*, *Prof. López is Marco*, but *Prof. López is not in the room*. These beliefs are also inconsistent. But here the inconsistency is not "built into" the concepts themselves.

Now, the incoherentist also thinks that it's a rule to avoid believing contradictions. So: a competent user of a concept is given a set of rules, all of which they cannot follow. Some (WRIGHT, 1975) take this impossibility as

a reductio of the incoherentist thesis. How can we play a game with rules that cannot all be obeyed?

Incoherentist respond: rules come in different strengths. Rules can be more or less strong, more-or-less central to the meaning of a concept, more-or-less negotiable. It's true that a competent user cannot follow all of the rules. Something must give. But a competent user will follow the rules as best they can, prioritizing the stronger, central, non-negotiable rules over the weaker, peripheral, negotiable ones. Presumably the rule for the soritical inference (R3) is weaker than the other rules. Hence, a competent speaker will ultimately resist their temptation to apply that rule in order to satisfy the others.

The incoherentist has an immediate explanation of the sorites susceptibility of vague concepts. We are inclined to make the paradoxical inferences because this inclination is required as a matter of conceptual competence! They have a similar explanation for the irremediable ignorance of middling cases. The inconsistency in the inference rules generates an instability in our judgments of middling cases. From one direction, we infer based on rule 1 and 3 that the middling case is one of baldness. From the other direction we infer based on rule 2 and an analogue of 3 that the case is one of non-baldness. Plausibly, in order to know p, we must have stable belief in p. So, we cannot know in middling cases. Incoherence generates ambivalence in middling cases, which precludes knowledge.

Suppose the incoherentist is right that vague terms are associated with inconsistent inference rules. What does that have to do with question closure? If anything, we should expect agents to feel a deep and persistent ambivalence towards middling cases, not question-closure!

The answer is that what we believe — and thus what inferences we draw — depends on which questions we see. Consider the following two examples discussed by Yalcin (originally given by Stalnaker):

> - William III of England believed that England could avoid war with France. Did he believe that England could avoid nuclear war with France? We want to say no...
> - The absent-minded detective believes the butler did it, but totally overlooks the possibility it was the chauffeur. Were he to consider this possibility, it might shake up his view. Does the detective believe that the chauffeur did not do it? We want to say no... (Yalcin, 2018, 34)

Both William III and the absent-minded detective have some beliefs without having an entailment of those beliefs. Compare these two with:

> - The foolish president believes that the US can avoid war with North Korea. But after asking himself whether or not the US can avoid nuclear

war with North Korea, he is left undecided about this further claim.

- The confused detective believes the butler did it. But after asking himself whether it was the chauffeur, he is left undecided about this further claim.

Like William III and the absent-minded detective, the foolish president and the confused detective also have some beliefs without having an entailment of those beliefs. But there is an obvious difference in their mental states. In the first pair of cases, the agents fail to see certain questions. In Yalcin's words:

> William III had, of course, no thoughts about matters nuclear *per se*....he lacks the concept NUCLEAR. ...his doxastic state is not defined on any resolution of logical space sensitive to questions of nuclear war. Again, it nowhere carves logical space according to whether nuclear war *per se* will, or will not, occur.
>
> The absent-minded detective...is capable of being sensitive to the relevant possibilities. He just isn't....His state of mind is defined on the question: *Is it the butler or not*? But it is undefined on the question: *Is it the butler, or the chauffeur, or someone else*? The possibility overlooked is inaccessible and backgrounded. (34)

We can think of questions as ways of carving up conceptual space, "resolutions" of conceptual space. But we don't always carve up conceptual space in all of the ways possible — there are certain questions that we may overlook, in the way that William III and the absent minded detective do. As Yalcin puts it, these are questions we are not sensitive to. As I will put it, these are questions we *fail to see*, questions that we are blind to.

Yalcin develops a detailed and powerful model of this feature of doxastic states. But for our purposes here, it suffices to simply note that there are two ways an agent might not draw an inference: they might be blind to the relevant questions (as with William III and the absent-minded detective) or they might see the relevant questions but simply fail to answer it in the way mandated by their other views (as with the foolish president and the confused detective).

With the notion of being blind to a question in mind, consider again the inference rules R1-R3. Those rules were of the form $x \approx y$. Flat-footedly, we might have interpreted such a rule as: if an agent believes x, then they should believe y. But now we can see that this interpretation smuggles in a requirement that the agent see the question whether y. It's best to recognize any such requirement as a separate rule. So let's separate these requirements as follows. We'll interpret $x \approx y$ as: if an agent believes x *and sees the question whether* y, then they should believe y. And we can add the separate rule-schema:

R4. The agent should see question Q.

This allows us to differentiate the mistake of the absent-minded detective — who failed to see a question — from the mistake of the confused detective — who saw the question but failed to draw the correct inference.

Of course, a competent user of the concept still cannot simultaneously follow all of these rules. Something has to give. Assuming R1 and R2 are non-negotiable, users can still resolve the conflict by breaking R3. But, on this way of individuating the rules, a new option for compromise emerges. Users might instead be like the absent-minded detective and be blind to the question of whether a middling scalp is bald. They might break an instance of R4.

Failing to see the relevant question, I claim, is the better resolution of the incoherence. The obligation encoded in R3 is stronger than our obligation to see a given question. So users take the former option and fail to see those questions. This is the best explanation of the data. If we refuse to see a question, we cannot take a questioning stance towards that question. So, question closure is to be expected. We cannot even be curious about p when we refuse to distinguish p in conceptual space — no more so than William III can be curious about nuclear wars. Once we blind ourselves to the distinction in conceptual space, there just doesn't seem to be any fact of the matter to ask after!

Call this the erotetic theory of vagueness. That theory supplements incoherentism with two claims. First: there is a mental act of refusing to see certain questions. That is, we can blind ourselves to those questions. Second: doing so is the best way to resolve the conflict in our incoherent rules. Doing so allows us to insulate ourselves from contradiction while fully satisfying the inferential obligations mandated by our concepts.

Above, we noted that there is something "conceptually confused" about someone who is curious who the last bald man in the sorites is — at least when they aren't being prodded by a philosopher. Part of conceptual competence is knowing which rules are stronger than others, and therefore which rules to give up in cases of conflict. A subject whose default state is to be curious about the last bald man is failing in that respect: they must be incorrectly prioritizing R3 over the relevant instances of R4. So, the subject is evincing conceptual confusion. In contrast, in order understand someone who is curious about (say) the temperature of the sun, we need simply to attribute certain strange curiosities to him — no conceptual confusion is required.

***

One immediate worry worth addressing arises when we consider how incoherentism (and our extension of it) applies to the case of OUGHT. Unlike the disjunctivist, I want to give a unified explanation for the vagueness-related features that we see across the descriptive and normative spheres. Those features included, for instance, sorites-susceptibility. But earlier, we noted (with Weatherson (2009) that there is nothing conceptually confused about someone — e.g. Prem, the flat-footed exceptionless deontologist — who does not find our original promise-vs-suffering series genuinely soritical. OUGHT is conceptually thin, in that there is no particular first-order normative theory that users of OUGHT are conceptually required to accept. This might lead us to be dubious that incoherentism can help us explain features like sorites susceptibility for OUGHT.

To see why we might be dubious, it's helpful to compare the case with BALD. It's part of the meaning of BALD that its application to a subject is determined by the number and distribution of hairs on that subject's head. That is to say, BALD has a conceptually determined parameter of application. Moreover, according to the incoherentist, competent users of BALD are required to (have a default disposition to) be tolerant to small changes in this parameter of application. That was how they explained the concept's sorites susceptibility and other vagueness-related phenomenon. However (à la Weatherson) OUGHT has no conceptually determined parameter of application. So, there is no parameter for which we might posit a tolerance rule. Thus, we might doubt that OUGHT is conceptually incoherent after all.

I have two replies to this worry.

First. Suppose the preceding worry is correct and OUGHT is not conceptually incoherent. So what? If we can't call the concept OUGHT vague, then so be it. (Perhaps we should call it "quasi-vague"?) What's important is whether or not we can still give a unified explanation of the vagueness-related phenomena that we observe with the concept. Importantly, whether or not we observe that phenomena in a given subject will depend on the first-order normative commitments that subject has. If, like Theresa, the subject believes that one ought to keep their promises iff doing so doesn't cause significant suffering, that subject will find the normative sorites series compelling and middling cases epistemically irremediable. More generally, subjects that base their ought-judgments on vague first-order normative theories will display the relevant phenomena. But then our explanation of the phenomena these subjects display can appeal to the fact that they accept these normative beliefs. Take Theresa. Using her normative beliefs, we can explain why she finds our original ought-series soritical. For two adjacent cases in our sorites, Theresa is conceptually mandated to be attracted to infer

from (i) *the first case doesn't cause significant suffering* to (ii) *the next case doesn't cause significant suffering*. And, given Therea's normative theory, she will infer from *you ought to keep your promise in the first case* to (i) and from (ii) to *you ought to keep your promise in the next case*. We can give a similar explanation for why Theresa can't come to know the exact point at which Regina should break her promise. Even if Theresa's normative theory is true, she cannot form a stable belief as to whether or not a middling case causes significant suffering. So, given Therea's normative theory, she also cannot form a stable belief as to whether Regina ought to keep her promise. (And because stable belief is a requirement for knowledge, she cannot know whether Regina ought to keep her promise.) We've thus given a unified explanation of the data, one that is parasitic on the incoherence of the descriptive concepts that appear in the subject's normative beliefs.

Here's my second reply. I agree that OUGHT has no conceptually mandated parameter of application. Nevertheless, I believe that OUGHT is indeed conceptually incoherent. The problem with the above argument is one of scope.[10] What we've noticed is that no *particular* parameter of application is such that users are conceptually mandated to treat OUGHT as tolerant with respect to *that* parameter. Nevertheless, it's plausibly conceptually mandated that we accept that there is some parameter of application or other for OUGHT.[11] And it may be conceptually mandated that *whatever* parameter one takes to ground applications of OUGHT one must treat OUGHT as tolerant with respect to that parameter.

We might remain skeptical of this suggestion. Is it really conceptually incoherent for someone to accept a precise cut-off for OUGHT? Recall, however, that we are, in the first instance, focused on the *decisional* OUGHT (see §2). It's plausible that users of the decisional ought are conceptually required to (have a default disposition to) treat the concept as abiding by certain "internalist" maxims. For instance, take two choice situations that are subjectively indistinguishable for the agent, in the sense that the difference between the choice situations is so small that they cannot tell that they are different (e.g. one choice situation involves an imperceptible amount of additional pain). Plausibly, agents are required to (have a default disposition to) move from *I decisionally ought to ϕ in the first case*, to *I decisionally ought to ϕ in the other case*. (Note the temptation to think that decisional ought facts depend only on our "internally accessible" states and are always knowable to us on that basis.) But, of course, subjective indistinguishability is not a transitive relation. So

[10]This is essentially the reply that Eklund gives to a similar worry raised by Weatherson with respect to GOOD. See WEATHERSON (2009, 86). (Weatherson remains unconvinced.)

[11]Cf. Zangwell's (2017; 2006) "Because Constraint".

this inference quickly gets us tangled in a sorites paradox — it is an inference that is ultimately incoherent. Of course, sustained philosophical reflection might lead us to resist our default disposition to accept such inconsistent inferences. But it's hard to deny its *initial* pull.

***

Plenty of the details of the erotetic theory remain to be filled in.[12] But I hope to have said enough to convince the reader that the theory, at least in broad strokes, is a promising explanation for question-closure.

I will now turn to the second half of our opening puzzle. In OUGHT sorites series, questions remain open for middling cases. Why is there persisting perplexity in such cases?

5 Step Two: Escapable and Inescapable Questions

Let's quickly recap. According to incoherentists, vagueness is the result of incoherent inferential rules inherent in our concepts. We cannot satisfy all of those rules while seeing all of the relevant questions. So, we face a choice. We must either break one of the inferential rules or blind ourselves to the relevant questions. I claimed that for typical vague questions, we have more reason to drop the relevant questions than to break the inferential rules. We should accept the preceding claim about our reasons because it's the best explanation of the data: it predicts the question-closure we observe.

On this approach, a strategy emerges for solving our previously intractable puzzle: we should focus on the differences between zetetic reasons in the descriptive versus normative spheres. That is, we should focus on the comparative strength of our reasons to see baldness-questions versus ought-questions. If we have more reason to see ought-questions than bald-questions, we can explain our puzzle.

In what follows, I will try to demonstrate the fruitfulness of this approach. I'll do this by presenting several arguments that the bar for refusing to see ought-questions is particularly high. The theme of these arguments will be that, unlike vague descriptive questions, vague ought-questions are *inescapable* for deliberative practical agents like us.

A warning before we proceed. My goal will not be to convince the reader of any one of these particular arguments. Rather, my goal is to convince the reader that the zetetic-framework is a productive one. Accordingly, I've

[12] I attempt to fill in some of those details in current work-in-progress.

chosen breadth over depth. I've chosen to highlight multiple provocative lines of thought instead of offering a plodding defense of one.

5.1 Practical Inescapability

There are at least two senses of "inescapable" for which one might plausibly argue that ought-questions are inescapable. The first picks up on a theme by Korsgaard and Sartre that choice is, in some sense, practically necessary.[13] As Korsgaard puts it:

> Human being are condemned to choice and action. Maybe you think you can avoid it, by resolutely standing still, refusing to act, refusing to move. But it's no use, for that will be something you have chosen to do, and then you will have acted after all. Choosing not to act makes not acting a kind of action, makes it something that you do. (Korsgaard, 2009, 1)

On one way of reading Korsgaard, she is claiming that one literally cannot decide not to act in a given choice situation. What might seem like a decision not to act in that situation is in fact a particular decision of how to act. For instance, it is a choice to act on the basis of your psychological states freed of normative governance — for short, we'll say that it's to choose to act *on a whim*. So, choosing not to act is literally something we cannot do.

We might reply: but we *do* sometimes choose not to act. After all, I can choose to get so blindingly drunk that my behavior later tonight will no longer count as action. Perhaps I will be *behaving* on a whim, but I am not *acting* on a whim.

That's an accurate observation. But I think Korsgaard's insight survives. It's true that I can now choose to get so blindingly drunk that my behavior later tonight will no longer count as an action. But the availability of this option doesn't release me from the responsibility of choice. When deciding now whether to get blindingly drunk, I must decide whether it's permissible for me to behave as a drunken fool tonight — even if such behavior doesn't count as an action. That's the sense in which we are condemned to choice. In other words:

K1. As rational agents, we decide not to act in a situation only if we decide that it's permissible to behave on a whim in that situation.

Plausibly, refusing to see an ought-question just is a way of deciding not to act in a situation. When we act, we decide what to do. And in order

[13] Thanks to Claire Kirwin and Zach Thornton for encouraging me to pursue this line of argument.

to decide what to do, we must decide what we ought and ought not do. Or, at the very least, we must be able to recognize various considerations regarding what we ought to do in that situation, and recognize them as such. But that's not something we can do if we've blinded ourselves to the question of what we ought to do in a given situation. Thus:

K2. Refusing to see the question of what one ought to do in a given situation is to decide not to act in that situation — to decide to behave on a whim.

Therefore:

K3. As rational agents, we refuse to see the question of what we ought to do in a given situation only if we decide that it's permissible to behave on a whim in that situation.

But, when we're considering a middling-ought case, we don't decide that it's permissible to behave on a whim. For instance, in the midst of Regina's promise-vs-suffering series, we don't know whether Regina ought to help the stranger. But if she ought to help the stranger, it's not permissible to behave on a whim. So, we don't decide that it's permissible to behave on a whim in her choice situation.

K4. As rational agents, we cannot and do not decide that it's permissible to behave on a whim in a middling situation.

Thus:

K5. As rational agents, we cannot and do not refuse to see the question of what we ought to do in a middling situation.

If we can't refuse to the see the question, we are stuck instead in a state of ambivalence — exactly the phenomena that we see.

There are, of course, lots of hidden assumptions in the above argument. For instance, I am assuming that in coming to a view about what *another* agent, e.g. Regina, ought to do in such-and-such a situation, I am also coming to a view about what *I* should do in that situation — what I ought to do were I in Regina's shoes (cf. Gibbard (2003)). Those are surely controversial claims. But instead of pausing to defend them, I will instead turn to a second sense in which ought-questions may be said to be inescapable.

5.2 Inquisitive Inescapability

Understood in this second sense, to say that ought-questions are inescapable is to say that they serve a necessary inquisitive purpose. Here's what I mean.

When faced with a typical vague question — say whether Harry is bald — the vagueness prevents us from coming to know the answer to that question. However, note that we can typically substitute a precise question which is equally well-suited for our inquisitive purposes. That is, there is typically a surrogate question which gives the inquirer at least as much information as they sought out with the vague question.

For example, suppose I am faced with the question of whether Harry is bald, where Harry is a borderline case of baldness. Unfortunately, because of the vagueness, I cannot come to know the answer to that question. Fortunately, I needn't worry: there are similar questions I can ask. For instance, I can ask whether Harry has at least $50k$ hairs on his head. Or I can ask how many hairs Harry has on his head. These are questions whose answers I *can* come to know, because they deploy (relatively) precise concepts.

I claim that the availability of these surrogate questions makes it less costly to refuse to see the original vague question. After all, we're no worse off given the inquisitive purposes we have in asking the question. And this is part of what explains why we're better off dropping the question than breaking a rule of inference.

Note, however, that for our thinnest ought-questions, reasonable surrogate questions are harder to come by. Suppose I ask what I decisionally ought to do in a given situation. You reply by introducing some precise notions — ought_1, ought_2... — and tell me that I ought_1 to do one thing, but ought_2 do something else. Is my inquisitive aim satisfied? No. That's because when I asked *what I ought to do*, my aim was to decide *what to do* in that situation. The aim is a practical rather than a theoretical one. And, I claim, the only rational way to satisfy my practical deliberative purpose is to form a view about what I decisionally ought to do. That's because there *are no* alternative questions that serve my original inquisitive purpose in asking the ought-question — no other questions that answer the practical question of what to do. Besides, *even if* these alternative questions were available in principle, it's hard to see why, from my current perspective, I *ought to let* my practical decisions be guided by them. After all, I am aiming to do what I ought to do, not what I "schmought" to do!

This gives us a view on yet another sense in which ought-questions (unlike bald-questions) are inescapable. And this helps explain why dropping ought-questions comes with a higher cost to our inquisitive aims than

dropping e.g. bald-questions.

In what follows, I'll try to state the foregoing argument with more rigor. I'll do so by developing the notion, introduced by Thomas Hofweber, of an *inescapable concept*. An inescapable concept, according to Hofweber, is a concept that "one cannot rationally, by one's own lights, replace...with another one for the purpose of inquiry" (2023, 8).[14] Extending Hofweber's notion, let's say that:

Inescapable Questions A *question* is *inescapable* when one cannot rationally, by one's own lights, replace it with another one for the purpose of inquiry.

First, I will argue that vague descriptive questions are typically *not* inescapable. Second, I will argue that vague ought-questions — at least for the thinnest ought concepts — are inescapable.

My argument for the first claim rests on two premises.

Premise 1: The inquisitive purpose of asking a typical descriptive question is to accurately represent, among the possible properties the world might have, which properties it in fact has. Equivalently, the purpose is to locate oneself in the space of metaphysically possible worlds. A question partitions the space of possible worlds, with each cell corresponding to an answer to the question. Upon coming to believe an answer to the question, we thereby locate ourselves in a particular cell.[15]

Premise 2: There is no "metaphysical vagueness": even though we can represent the properties of the world using vague concepts and vague propositions, the properties themselves are not vague. At a minimum, this means that there are no two metaphysically possible worlds that make all of the same precise propositions true, but make different vague propositions true (cf. Bacon (2018, §15)). That is, there are not two distinct metaphysically possible worlds: one where Harry has exactly n hairs (arranged thus-and-so) and is bald, and one where Harry has exactly n hairs (arranged thus-and-so)

[14]Eklund (2015) discusses a similar (although slightly broader) notion of a *conceptual fixed point*.

[15]Even questions expressed with a posteriori necessities can be understood in this way. For instance, if an early chemist asks "is water H_2O?", they are presumably still comparing worlds: namely those where the natural kind that fills our lakes and oceans, comes out of our taps, etc. is H_2O, and worlds where a different natural kind (e.g. XYZ) or no natural kind plays that role. That's a division in the space of possible worlds. Admittedly, it's much harder to apply the foregoing world-property model of inquisitive purpose to conceptual questions – questions like *are all bachelors unmarried?*. But I expect that we can cash out the inquisitive purpose of asking conceptual questions as derivative on the purpose of asking non-conceptual questions.

and is not bald. Having exactly n hairs necessarily makes Harry bald, or it necessarily makes him not bald — even if, for middling n, it is vague which.

With these two premises, it's a short step to argue that questions like *is Harry bald?* are not inescapable. In asking this question, as with other descriptive questions, my purpose is to come to locate myself in the space of possible worlds — to figure out which of the possible world-properties are instantiated. Asking this question is meant to further that goal by dividing the space of possible worlds into two cells: worlds where Harry is bald and those where he is not bald. And answering this question helps me decide which cell I am located in, thereby helping me narrow in on a set of worlds as candidates for actuality.

Unfortunately, when Harry is a middling case of baldness, I am unable to answer this question. My ignorance, however, is due to a defect in the question rather my ability to detect world-properties. That is, it's vague which worlds are in which cell — it's vague which worlds are ones where Harry is bald and ones where he is not bald. I can better accomplish my inquisitive purpose by asking the less defective question *how many hairs does Harry have?*. This question divides up the space of possible worlds into many cells, grouping worlds where Harry has the same number of hairs on his head. This question is strictly more fine-grained than my previous vague-question — each cell in the previous question will be a union of cells in this question. And it might take more cognitive work to consider this more fine-grained question. But doing so serves my inquisitive ends better than my original question. Whatever properties of the world I sought to represent in asking the first question, I can come to represent it by asking this "higher resolution" surrogate question. So, that original question is not inescapable.

Turn now to ought-questions. I'll argue that questions for our thinnest ought-concepts (and in particular the decisional OUGHT) are inescapable. My arguments rely on the following assumption: when we inquire with ought-questions, the purpose is not to locate ourselves among the space of possible worlds. Rather the purpose is distinctly practical: we are trying to decide *what to do*. More carefully: the inquisitive purpose of the question of what we ought to do in a given situation is to rationally guide our action in that situation.

Take this premise and suppose for reductio that ought-questions are escapable. Then there is some distinct question of the form *what schmought I do in such and such a situation?* that, by my own lights, does a better job rationally guiding my action. I will give several arguments that there are no such schmought-questions, and thus that the only question we ought to let

rationally guide our actions are the ought-questions themselves.

My arguments come in two varieties. The first *strong* variety ends with the conclusion that there are no alternative schmought-questions available to guide my actions. That is, it ends with the conclusion that there are no concepts distinct from OUGHT that play the same practical role as OUGHT. A second *weaker* variant concludes that, even if alternative schmought-questions are available to guide choice, it would be irrational to let these alternative questions steer my action.[16]

5.2.1 Strong Inquisitive Inescapability

Turning to the first variety, here's an argument that exploits linguistic intuitions like those in moral twin earth (HORGAN AND TIMMONS, 1992). Suppose that we believe that we ought to act in accordance with causal decision theory. So, when faced with a Newcomb problem, we decide that we ought to choose both boxes. Letting this rationally guide our actions, we form the intention to choose both boxes in those cases. We then meet a tribe of speakers who claim to understand our concept OUGHT but to act on the basis of another concept SCHMOUGHT. They say:

> We agree that one ought to act in accordance with causal decision theory — that we ought to take both boxes. But we act on the basis of what we schmought to do. And we schmought to act in accordance with evidential decision theory. So, in Newcomb cases, we schmought to take one box. On that basis, I'll go ahead and just take the one.

Intuitively, these tribe members are misdescribing their own thoughts. We disagree with them: they think one ought to one-box and we think one ought to two-box. What they are calling the concept "SCHMOUGHT" is just our concept OUGHT. And what they are calling the concept "OUGHT" is just the descriptive concept RECOMMENDED BY CAUSAL DECISION THEORY.

This makes plausible that the meaning of our OUGHT concept is fully-determined by the inferential role that concept plays in practical deliberation, and, specifically, in guiding our choice. The fact that we base our decision to do x in circumstance C on a judgment about x in C is what makes that judgment the judgment that we ought to do x in C. There is no alternative question that better serves that purpose because there are no alternative questions that *can* serve that purpose — the purpose defines the question! So, ought-questions are inescapable.

[16]The arguments for "weak" inquisitive inescapability more closely mirror the sorts of arguments given in HOFWEBER (2023).

Really, I'm taking this thought experiment to support two claims. First: I'm claiming that OUGHT is associated with an inferential role that is practical. For instance, OUGHT might be associated with something like the following inference: We move from the belief I OUGHT TO DO x IN C to the decision IF C, DO x!. Second: I'm claiming that this practical-inferential role fixes the extension of the concepts that play it – that OUGHT has what Eklund (2017, 10) calls "referential normativity". Taken together: there isn't an alternative concept that both plays the OUGHT-role and yet applies to different options.[17]

On the first claim: Those with strong motivational-externalist commitments will admittedly not find this line of thought convincing. But for those who are undecided, note that positing a practical-inferential role helps us explain (or rather, explain *away*) part of The Puzzle of Bouletic Exceptionalism. Recall that puzzle: there was a general ban on intrinsically desiring the vague, but yet it seemed perfectly rational to intrinsically desire that we do what we ought to do. That's puzzling — why the exception? However, if ought-thoughts figure in practical inferences, we can resist the seemingly obvious claim that we desire to do what we ought to do. It might seem like we have a desire to do what we ought to do because we are highly motivated to do what we think we ought to do. But, we can now see that positing such a desire is not necessary. Our ought judgments motivate us to action via a direct practical inference. No "fetishistic" normative desire is required to mediate.

On the second claim: there are alternative arguments in the literature that further support the claim that the practical-inferential role of a concept like OUGHT determines its extension. For instance, Eklund (2017; 2012) argues that denying this assumption opens us up to the practical question of whether to do what we ought to do or, instead, what we schmought to do. And he argues that — given that each concept picks out its own distinct property — both answers would be "objectively backed by reality". That, he takes it, is antithetical to the spirit of realism.[18] If we treat this conclusion as a reductio, we have additional reason to think that there aren't alternative schmought-questions to begin with, that the practical-inferential role of OUGHT does determine its extension.

[17]For another argument along these lines, see VERMAIRE (2021). There, Vermaire presents a careful and compelling argument that, whether or not the referential normativity thesis is true or false, the practical role of ought precludes an agent from rationally acknowledging that there are different ought/schmought questions. Eklund responds in his (2023).

[18]See especially Chs. 1 & 2 of EKLUND (2017).

5.2.2 Weak Inquisitive Inescapability

Here's an argument (closely following Hofweber (2023, 11)) of the second variety. To say that the precise schmought-question, by my own lights, does a better job rationally guiding my action is just to say that I ought, by my own lights, to change the basis of my acting from the ought-question to the precise schmought-question. But, I ought, by my own lights, to change that basis iff I know that by doing so I'm more likely to do what I ought to do. But this can't be the case! If schmought is some precise version of ought, then I can't know whether or not it has the same extension as my vague concept ought. So, I cannot know if guiding my action by schmought thoughts is more likely to lead me to do what I ought to do. So ought-questions are inescapable.

Of course, were I to make the switch, I would be more likely to do what I schmought to do. Indeed, once the switch is made, I would endorse it: I would correctly conclude that I schmouldn't base my actions on what I ought to do. And I can know all of this from my current perspective. But that's not relevant to my present evaluation. When I evaluate which questions I *ought* to ask when deciding what to do, I am using my very concept of ought. Ought-questions thereby form a sort of zetetic fixed-point.

6 Conclusion

We began by noting that the middling cases in sorites series for ordinary descriptive concepts exhibit question-closure. And we noted that this is not the case for thin normative concepts, even though they exhibit the other vagueness-related phenomenon. In the descriptive case, we explained question-closure as the result of our refusal to see certain questions, which in turn was explained as the least costly way to insulate ourselves from incoherence. For ought-questions, however, the cost of blinding ourselves is much higher. Those questions are inescapable.

This explanation has some obvious dialectical upshots. For instance, it yields support for incoherentism. And in so far as we don't need a disjunctive strategy to explain the "hidden fact of the matter" phenomenology, we undercut the motivation for radical metaphysical posits like a *sui generis* parity relation and hidden joints in moral reality.

Perhaps most importantly, if my theory is right, it supports a sort of quietism. Here's what I mean. There has been lots of work asking how we should act in various cases of indeterminacy. How should we act when it's vague which action is best for us? How should we act when it's vague

which action is morally right?[19] Traditional decision rules (e.g. we should act to maximize expected utility, or we should do whatever is morally right) tell us that it's *vague* what we ought to do. We are, in my terminology, in a quasi-borderline case. That is an uncomfortable resting point for us philosophers. It feels as though there is some "hidden fact of the matter" about what we ought to do, one that we can be anxious to discover. As a result, we feel compelled to devise elaborate modifications of the traditional theories, hoping to give us some guidance in cases of decision making under indeterminacy. But, if I'm right, the perplexity is not a sign that further guidance is possible. Despite the phenomenological difference, there is no hidden ought-fact which we philosophers can hope to discover, any more than we can discover whether the borderline-bald are in fact bald. Instead, we're better off learning to live with the persisting perplexity.

References

Andrew Bacon. *Vagueness and Thought*. Oxford University Press, 2018.

David Chalmers. Verbal disputes. *Philosophical Review*, 120(4):515–566, 2011.

Ruth Chang. The possibility of parity. *Ethics*, 112(4):659–688, July 2002.

Cristian Constantinescu. Value incomparability and indeterminacy. *Ethical Theory and Moral Practice*, 15(1):57–70, February 2012.

Cian Dorr, Jake Nebel, and Jake Zuehl. The case for comparability. *Noûs*, 57: 414–453, 2023.

Tom Dougherty. Vague value. *Philosophy and Phenomenological Research*, pages 352–372, 2013.

Michael Dummett. Wang's paradox. *Synthese*, 30:301–24, 1975.

Matti Eklund. Inconsistent languages. *Philosophy and Phenomenological Research*, 64(2):251–275, March 2002.

Matti Eklund. What vagueness consists in. *Philosophical Studies*, 125:27–60, 2005.

Matti Eklund. Alternative normative concepts. *Analytic Philosophy*, 2(53): 139–157, June 2012.

[19] See, for instance, Williams (2017).

Matti Eklund. Intuitions, conceptual engineering, and conceptual fixed points. In Chris Daly, editor, *The Palgrave Handbook of Philosophical Methods*, pages 363–85. Palgrave Macmillan, 2015.

Matti Eklund. *Choosing Normative Concepts*. Oxford University Press, 2017.

Matti Eklund. Incoherentism and the sorites paradox. In Sergi Oms and Elia Zardini, editors, *The Sorites Paradox*, pages 78–94. Cambridge University Press, 2019.

Matti Eklund. Schmoughts for naught? Reply to Vermaire. *The Journal of Philosophy*, 120(7):392–398, July 2023. doi: 10.5840/jphil2023120721.

Luke Elson. Incommensurability as vagueness: a burden-shifting argument. *Theoria*, 83:341–363, 2017.

Hartry Field. Indeterminacy, degree of belief, and excluded middle. *Noûs*, 34:1–30, 2000.

Hartry Field. This magic moment: Horwich on the boundaries of vague terms. In Richard Dietz and Sebastiano Moruzzi, editors, *Cuts and Clouds*, chapter 11, pages 200–208. Oxford University Press, 2009.

Jane Friedman. Inquiry and belief. *Noûs*, 53(2):296–315, 2019. doi: 10.1111/nous.12222.

Allan Gibbard. *Wise Choices, Apt Feelings*. Harvard University Press, 1990.

Allan Gibbard. *Thinking How to Live*. Harvard University Press, Cambridge, Massachusetts, 2003.

Allan Gibbard. Truth and correct belief. *Philosophical Issues*, 15:338–350, 2005.

Thomas Hofweber. Inescapable concepts. *Australasian Journal of Philosophy*, 2023. doi: 10.1080/00048402.2023.2172593. Advance online publication.

Terry Horgan and Mark Timmons. Troubles for new wave moral semantics: The open question argument revived. *Philosophical Papers*, 21:153–75, 1992.

Christine M. Korsgaard. *Self-Constitution*. Oxford University Press, 2009.

David Lewis. New work for a theory of universals. *Australasian Journal of Philosophy*, 61(4):343–377, December 1983.

S. Matthew Liao. Threshold deontology: Lessons from vagueness. Talk given at "Vagueness and Ethics" Workshop at Uppsala University on May 4, 2023.

Stephen Schiffer. *The Things We Mean*. Oxford University Press, 2003.

Miriam Schoenfield. Moral vagueness is ontic vagueness. *Ethics*, 126:257–282, January 2016.

Jack Spencer. Intrinsically desiring the vague. *Proceedings of the Aristotelian Society*, 122(Part 2):165–194, 2022.

Matthew Vermaire. Against schmought. *The Journal of Philosophy*, 118(9): 465–485, September 2021.

Ryan Wasserman. Indeterminacy, ignorance, and the possibility of parity. *Philosophical Perspectives*, 18:391–403, 2004.

Brian Weatherson. Vagueness as indeterminacy. In Richard Dietz and Sebastiano Moruzzi, editors, *Cuts and Clouds*, chapter 4, pages 77–90. Oxford University Press, 2009.

J.R.G. Williams. Indeterminacy, angst, and conflicting values. *Ratio*, 29(4): 412–433, 2016.

J.R.G. Williams. Indeterminate oughts. *Ethics*, 127(3):645–673, 2017.

Crispin Wright. On the coherence of vague predicates. *Synthese*, 30:325–65, 1975.

Seth Yalcin. Belief as question-sensitive. *Philosophy and Phenomenological Research*, 97(1):23–47, July 2018. doi: 10.1111/phpr.12330.

Nick Zangwill. Moral epistemology and the because constraint. In James Dreier, editor, *Contemporary Debates in Moral Theory*, pages 263–281. Blackwell, 2006.

Nick Zangwill. Moral dependence and natural properties. *The Aristotelian Society*, 91:221–243, 2017. doi: 10.1093/arisup/akx007.