Philosophia (2024) 52:161–186
https://doi.org/10.1007/s11406-024-00716-6

Williamson's Epistemicism and Properties Accounts of Predicates

Paul Teller[1]

Received: 3 January 2024 / Accepted: 3 January 2024 / Published online: 26 February 2024

Abstract
If the semantic values of predicates are, as Williamson assumes (*Philsophical Perspectives, 13*, 505–517, 1999, 509) properties in the intensional sense, then epistemicism is immediate. Epistemicism fails, so also this properties account of predicates. I deploy examination of Williamson's account as a foil against properties as semantic values, showing that his two positive arguments for bivalence fail, as do his efforts to rescue epistemicism from obvious problems. In Part II I argue that, despite the properties account's problems, it has an important role to play in compositional semantics. We may separate the problem of how smallest parts of language get attached to the world from the problem of how those parts compose to form complex semantic values. For the latter problem we idealize and treat the smallest semantic values as properties (and referents). So doing functions to put to one side how the smallest parts get worldly attachment, a problem that would just get in the way of understanding composition. Attachment to the world must be studied separately, and I review some of the options. As a bonus we see why the requirement of higher order vagueness is an artifact of taking properties as semantic values literally instead of as a simplifying idealization.

Keywords Vagueness · Properties · Williamson · Semantics · Use accounts of meaning

1 Introduction

So many have expressed incredulity towards Williamson's (1992, 1994a) epistemic account of vagueness, why another examination? The interest is the simple, attractive, and widely held presumption about the meaning of predicates, that predicates express properties (always understood one or many place).

We can't take the meaning of a predicate to be its extension - then a predicate could not have different extensions in different (actual, counterfactual, and future) contexts,

✉ Paul Teller
prteller@ucdavis.edu

[1] Department of philosophy, University of California at Davis, Davis, CA, USA

and there could not be coextensional predicates with different meanings. A predicate must have some intermediary that will function to pick out its variable extension. Such intermediaries could be "properties" "characteristics" "truth conditions", "rules of application"…;[1] there are many options. My target will be accounts that take such intermediaries to determine exact extensions for predicates, that is extensions that are sets or like sets in that for every entity[2] either the entity is in the extension or is in the counter-extension. Such an intermediary could be a property in the minimal sense of an intension, that is a function from possible worlds (or sufficiently detailed circumstances) to sets of entities in that world, understood as the property's extension in that world. Or one might take properties to be some sort of ontologically more robust category, as long as each member of that category determines a property in the minimal intensional sense.[3] It will be the exact extensions that engender the problems, so by 'property' readers should understand the minimal intensional notion of a function from worlds to sets, or any stronger notion that entails this minimal notion. By "a properties account of predicates" I will understand any account that takes the meaning of predicates to be properties understood in the intensional sense just explained. [4]

Such a properties account immediately entails bivalence: If the meaning of a predicate, *P*,[5] is a property in this intensional sense, then in each world[6] the predicate has a set valued extension. But then in each world, for any object, *a*, either *a* is in the extension or it is in the counter-extension. Since the extension of a predicate, *P*, is the set of things in which the predication, *Pa*, is true and its counter-extension is the set of things in which ¬*Pa* is true it follows that either *Pa* is true (*a* is in *P*'s extension) or ¬*Pa* is true (*a* is in *P*'s counterextension). This is bivalence in the form, for any predication, *Pa*, either *Pa* is true or ¬*Pa* is true. If we follow Williamson (see below) in defining the word 'false' as *Pa* counts as false just in case ¬*Pa* is true we have bivalence in the form that every (contentful) predication is true or false.[7]

[1] I will use single quotation marks to turn an expression into a name of that expression and double quotes for all other uses such as shudder quotes, reporting what has been said, etc.

[2] I will follow Williamson and this literature generally and ignore complications arising from sortals, so that everything will count as a candidate. Bringing in sortals would needlessly complicate exposition while making no difference in the end to what will be at stake.

[3] For example, one might want to distinguish necessarily equivalent properties such as that of being triangular and being trilateral.

[4] Could vagueness be in the properties themselves so that properties may not have exact extensions? There is a growing literature exploring the idea of "worldly indeterminacy", but I am sketpical that such accounts would have application to the vagueness of predicates. Worldly indeterminacy of properties would, some exceptional cases aside, be things in the world, independent of human vagaries; while vagueness in human languages is shot through with considerations tied to the accidents of human language development and use. Detailed comparison with the present account will have to wait for another occasion.

[5] I will use uppercase italic '*P*' as a variable over predicates; lower case italic '*a*' as a variable over referring expressions, and '*Pa*' as a variable over predications, that I understand as a combining of predicate, *P*, and referring expression, *a*, to form a sentence that can be used to make an assertion.

[6] In general, different extensions in different worlds. For the rest of the paragraph everything is relativized to a world.

[7] This argument uses classical logic. For many non-classical logics bivalence fails. My approach in this paper is to critically examine Williamson's view assuming the classical logic to which he is committed and then to peruse the resulting approach to vagueness again assuming classical logic. Comparison with approaches with other logics requires first developing the present approach. Examination of the relative strengths and weaknesses of the various approaches must wait for another occasion.

In Part I I deploy critique of Williamson on vagueness as an extended argument against any properties account of predicates.[8] Williamson offers two positive arguments for bivalence. I show that the first argument fails and that the second ultimately presupposes the properties account. Epistemicism suffers insuperable difficulties, which Williamson's detailed exposition attempts to address. I extensively supplement arguments already in the literature that, together, show that Williamson's efforts fail. Any properties account immediately entails the bivalence of epistemicism.[9] Having shown that epistemicism's bivalence is indefensible, we conclude that we are in need of one, or more, alternatives to any properties account of the meaning of predicates.

Part II discusses prospects for alternatives. I first provide a way of thinking of the role of properties accounts that shows, despite the difficulties, how they can usefully function in semantic theory. Thinking of a properties account's role in the suggested way also provides a reconciliation between compositional semantics and so-called "use" accounts of meaning. I briefly explain how predicates can function in language even though they have no extensions, followed by a short survey of some of the ways, mostly familiar, in which the proposed approach to predicates might be filled out. As a bonus we see that higher order vagueness is not forced.

2 Part I: Critical Examination of Williamson's Epistemicism

2.1 Thumbnail Sketch of Williamson's Account

According to Williamson vague statements[10] are all bivalent, that is either true or false, very much including unclear borderline cases, both actual and counterfactual. Removing grains of sand from a heap, there is a last one which, when removed, leaves no heap behind. For each individual, there is a fact of the matter whether or not they are bald. Vagueness is not denied. Rather it is claimed to be an epistemic, not a semantic phenomenon. In unclear cases there is always a truth of the matter, but it is out of human reach to know just what that truth is. Williamson's epistemicism is simply bivalence, elaborated with the interpretive gloss that vagueness is not

[8] Elsewhere (2017) I have provided entirely different arguments against any properties account, (though not under that name). In "Vague so Untrue" (2007) Braun and Sider also reject properties accounts, providing (p. 135) a one sentence summary of the kinds of considerations that I develop in the article just mentioned. Ludwig and Ray (2002, to appear), provide yet another argument.

[9] In his (2011) Elkund reviews various ways in which authors have tried to defend bivalence in some way that they do not count as epistemicism, concluding (p. 358) that "I see the attempts to defend bivalence in a non-Williamsonian way as constituting one major recent trend." Epistemicism as understood in the present paper is bivalence of statements (Williamson: utterances that say that something is the case) plus the interpretive gloss. This is clearly Williamson's understanding (1994a, pp. 185 ff.) In any case the issue is moot: The arguments in this paper against epistemicism are all arguments against the bivalence of properties accounts. If some authors understand the term 'epistemicism' in other ways that isn't relevant to the present arguments.

[10] Williamson uses 'utterances' explicitly restricted to occasions when someone uses an utterance "to say that something is the case." (1994a, p. 187) I will use 'statement' and 'utterance', thus restricted, interchangeably, 'utterance' being the choice when discussing passages from Williamson where he uses that term.

a semantic but an epistemic phenomenon. Once bivalence is accepted, I take the interpretive gloss to go without saying.

Williamson presents two positive arguments for bivalence. Insisting on bivalence, Williamson claims, is the only approach to vagueness that does not require that we give up some part of classical logic. He then submits a proof that assumption of any counterexample to bivalence leads to a contradiction. The rest of Williamson's material attempts to counter arguments against epistemicism, and especially to counter epistemicism's manifest implausibility. A major part of Williamson's response to the implausibility of epistemicism is margin for error principles: "[M]argin for error principles explain both the ignorance postulated by the epistemic view and the apparent intuitions that run counter to that view." (1994a, 234) There are several extensive critical examinations of Williamson's appeal to margin of error principles.[11] As I take these critiques to be sufficient, I will focus on his prior, and in some ways more basic, efforts to accommodate epistemicism's manifest implausibility.

2.2 Williamson's Arguments for Bivalence

Williamson avers that

> If one abandons bivalence for vague utterances, one pays a high price. One can no longer apply classical truth-conditional semantics to them, and probably not even classical logic. Yet classical semantics and logic are vastly superior to the alternatives in simplicity, power, past success, and integration with theories in other domains. (1994a, 186)

Williamson takes this consideration to be far from conclusive, but still, as above, "a high price". (1994a, 186) [12]

Williamson has presented a false dilemma: One must take classical logic to apply for all or for no vague utterances. But when one can safely take the case at hand to be clear, bivalence can safely be assumed and there is no need for any alternative to classical logic. As Schiffer puts it, "In most cases we can harmlessly assume the premises in an argument have truth-values and then apply classical logic with assurance of truth preservation." (1999, 501) On the other hand, where one needs to treat a case as unclear, no truth values are accessible to us. That is what being an unclear case comes to. For such cases, classical logic will generally have no application whether or not there are inaccessible truth values. So there is no need to appeal to epistemicism to save classical logic.

Some will find this short repost a bit slick: What counts, they will say is whether an utterance *has* a truth value, not whether having a truth value can be safely or harmlessly assumed. But the things that really can be safely or harmlessly assumed

[11] Keefe (2000, 64–70), Ray (2004), Mahtani (2004 and 2008), Machina and Deutsch (2002, 35–45), Wright (1995, 148–152).

[12] Several comments in the literature are along the lines of what immediately follows: Schiffer (1999, 501), McGee and McLaughlin (1998, 225), Machina and Deutsch (2002, 29) Ludwig and Ray (2002, 429–431; 2017, Preprint, 6).

will *usually be* as assumed. We still get the conclusion that classical logic justifiably, and most often correctly, applies in the clear but possibly not in the unclear cases.

What may impress readers more is Williamson's valid proof of bivalence: Using classical logic, the Tarski biconditionals, and truth and falsity predicates in the object language (that I will write as 'true(s)' and 'false(s)'), Williamson proves true(s) v false(s).

Williamson (1992, 145-6 and 1994a, 187-9) formulates the argument as a reductio. But the crucial steps give an immediate direct proof. Using the Tarski biconditionals as premises

(T) true(s) <-> s
(F) false(s) <-> $-s$

and starting with the tautology

(1) s v $-s$

(T) and (F) license the substitutions of their left-hand sides for their right-hand sides, giving

(2) true(s) v false(s)

QED!

Should this proof move a bivalence skeptic? To answer this question most clearly it will be helpful to have a specific bivalence skeptic in mind. I will use supervaluationism as the exemplar, though much of what I say should go for many other bivalence skeptics. Here is a short summary of supervaluationism. (Fine, 1975; Keefe, 2000 ch. 7).

Supervaluationism considers all the admissible ways in which a vague term such as 'tall' could be made precise. These are called 'admissible precisifications'. Henceforth the 'admissible' will be taken for granted. On each precisification any relevant statement is taken to be classically true or false. A statement is said to be super-true just in case it is true on all precisifications, super-false if false on all precisifications. Supervaluationists then identify truth with super-truth, falsity with super-falsity. On the supervaluationist account, failures of bivalence are the cases in which a statement is (classically) true on some precisifications, (classically) false on others. Finally, for a given sentence, s, if s is super-true, it counts as super-true at all its precisifications, and similarly for super-false.

The identification of truth with supertruth is contentious- I will discuss this issue below. For the moment when what is in question is truth as interpreted by supervaluationism I will write '(super)true' and '(super)false'. When what is in question is an interpretation of truth that satisfies the Tarski biconditionals, I will write '(classically) true' and '(classically)false'. When the difference doesn't matter, I will write simply 'true' or 'false'. [13]

[13] At (1994a, 162–163) Williamson inveighs against (super)truth, but the argument there is to insist, several times over, that real truth must satisfy the Tarski biconditionals. Thus this question devolves upon Williamson's arguments for the biconditionals.

Supervaluationists reject Williamson's argument because they reject (T) and (F).[14] Consider some *s* that is (classically)true at some precisifications, (classically)false on others, so that *s* is neither (super)true nor (super)false. So, according to supervaluationists, 'true(*s*)' is (super)false. For this *s*, supervaluationists will count the right and left hand sides of (T) not to match in truth value. The same problem applies to (F).

This supervaluationst objection to Williamson's argument can be met by avoiding appeal to (T) and (F). As Williamson himself notes (1994a, 162-3. See also Keefe's discussion, 2000, 214–217) instead of (T) and (F) we can adopt four new rules of inference. Since on a supervaluationist account

$$\text{true}(s) \vDash s,$$
$$s \vDash \text{true}(s)$$

$$\text{false}(s) \vDash \neg s,$$
$$\neg s \vDash \text{false}(s)$$

are valid, supervaluationists can adopt the new rules of inference,[15]

(T*) $\text{true}(s) \vdash s$, $s \vdash \text{true}(s)$, $\text{false}(s) \vdash \neg s$, $\neg s \vdash \text{false}(s)$[15]

Williamson rejects T* as an adequate replacement for (T) and (F). (1994a, 162-3, Williamson) But what is at issue here is T*'s validity on a supervaluationists' account, which Williamson does not dispute. Using T* we argue:

1)	$s \vee \neg s$	
2)	s	Assumption for vE
3)	$\text{true}(s)$	2, T*
4)	$\text{true}(s) \vee \text{false}(s)$	3, vI
5)	$\neg s$	Assumption for vE
6)	$\text{false}(s)$	5, T*
7)	$\text{true}(s) \vee \text{false}(s)$	6, vI
8)	$\text{true}(s) \vee \text{false}(s)$	1, 2-4, 5-7, vE

Line 1 is a classical tautology, endorsed by supervaluationists. So it would appear that we have proved bivalence, whether the object language predicate 'true' is interpreted as (classic)truth or as (super)truth!

[14] Supervaluationists also reject truth functionality on which Williamson's argument depends. See Keefe and Smith (1999, 27) for a simple counterexample.

[15] 'T*' is Keefe's notation – see her (2000, 214) for argument and other details. Keefe proposes T* instead of (T) and (F) above not as a problem for Williamson's argument for bivalence but in the context of rejecting Williamson's claim that super-truth "is not disquotational." (1994a, 162) McGee and McLaughlin (1998, 224) claim that if bivalence isn't presupposed, T* (their T and F introduction) are not valid but they give no argument. What they do argue is that, where bivalence fails, no "truth preserving rule of inference can be validly employed within conditional proofs." As far as I can see, their argument begs the question, but should they be right, so much the worse for the attempted rescue below. Richard (2000) also provides considerations relevant to this discussion.

But supervaluationists have a new objection to this argument. Classical logic holds for supervaluationists, but only when the object language does not have the truth predicate.[16] When the object language includes the truth predicate, interpreted as (super)truth, v elimination (argument by cases) fails (as Williamson himself notes (1994a, 152). If *s* is neither (super)true nor (super)false, *s* v ¬*s* is still (super)true, since at each precisification one of, *s* or ¬*s* is true; but for such an *s*, (super)true(*s*) v (super)false(*s*) fails.

Where do we stand? If truth is interpreted classically the argument, and similar ones, are valid; and bivalence has been proven. If truth is interpreted as super-truth the argument is invalid. So Williamson has a case for bivalence if, and only if, truth is to be interpreted classically. Any more detailed account of classical truth will serve Williamson if, and only if, it is a notion that will support the Tarski biconditionals.

Williamson in fact acknowledges supervaluationism and its super-truth as at least an apparent counterexample to his argument. Williamson's argument for bivalence also appears in his (1992, 145–146). In note 7, p. 148 he writes: "The supervaluational treatment of vagueness…may seem an obvious counterexample to the argument." Williamson counters that: "Where the present approach differs is in its claim that the ordinary notion of truth is subject to the Tarskian schema and is therefore not to be defined [as super-truth]." (1992, 148) In other words, Williamson acknowledges that bivalence fails for (super)truth, but, (super)truth isn't *truth*. A similar pattern of discussion occurs at (1994a, 162–163) where Williamson inveighs against (super)truth insisting, several times over, that real truth must satisfy the Tarski biconditionals. So everything turns on the Tarski biconditionals and Williamson's argument for them.

2.3 Williamson's Argument for the Tarski Biconditionals and the Properties Account of Predicates

Here is Williamson's argument for the Tarski biconditionals:[17]

> The rationale for the disquotational character of truth is simple. Given that an utterance says that TW is thin, what it takes for it to be true is just for TW to be thin, and what it takes for it to be false is for TW not to be thin. No more and no less is required. (1994a, 190)[18]

Williamson considers the repost:

> It might be replied that if u says that P and is neither true nor false, then 'u is true' is false while P is neither true nor false, so that the two sides of [(T)] do not match in semantic value….The trouble with this objection is that it does

[16] Or their D[efinitely] operator, that functions very like the truth predicate.

[17] Torrago (1998, 638) explores the further option of taking the truth predicate itself to be vague, which would save the Tarski biconditionals in a way what won't help epistemicists. This option can be developed by reworking the concept of truth along the lines of "true enough". See Teller (2017) and Elgin (2017)

[18] Williamson repeats this argument (1997b, 217).

> nothing to meet the rationale for [the Tarski biconditionals]. It gives no hint, when u says that TW is thin, of any way in which u could fail to be true, other than by TW failing to be thin,…." (1994a, 190).
>
> What is this argument? I take Williamson to be applying a
>
> Default Principle: If a sentence, u, (assumed to be contentful) fails to be true, then, by default, it counts as false.

In the example in which u says that TW is thin, if u fails to be true then TW fails to be thin, in which case, by default, u counts as false. As for the repost, Williamson's "does nothing to meet the rationale…" seeks to shift the burden of proof: It's not good enough simply to comment that u could be neither true nor false. The Tarski biconditionals command such prima facie plausibility that to reject them one must, at the very least, show how they could fail. I will show that the default principle begs the question. Part II address the burden of proof by illuminating what is involved in a contentfull statement being neither truth nor false and by exploring a range of accounts that provide substance to the failure of the Tarski biconditionals.[19]

Let's examine a passage where Williamson explicitly spells out application of his default principle:

> To determine which property 'bald' refers to, the reference-determining factors must determine of each thing x, time t and possible world w whether x at t in w is to have the property, in other words, whether the ordered triple (x, t, w) is to belong to the intension of 'bald'. Nothing more is needed… All sides agree that whatever facts there are about the reference of 'bald' are determined by the reference-determining factors (such as use and the environment); the disagreement concerns what facts there are to be determined. Thus if there are not enough facts about use and the environment to determine (x, t, w) to belong to the intension of 'bald', then that very shortfall is enough to determine it not to belong, and is itself determined by the facts about use and the environment. Reference can go by default. The worry that there might not be enough facts about use and the environment to do the determining in every case is misconceived. (1999, 509)

In other places Williamson more briefly appeals to his default principle, for example (1994a, 208, 213–214; 1997b, 224–227).

Many readers will find, as I did, a lot of plausibility in the default principle. But consider: To apply the default principle there must already be, in a given possible world or circumstance, an extension in opposition to which the default principle applies. (In the immediately following "In a given possible world or circumstance" will be understood.) To apply the default principle to an x requires that x *not* be one such that "the reference-determining factors" have determined it to be in the extension of the predicate in question. But if an

[19] Sainsbury (1999, 259) describes what amounts to the default principle as central to what he calls the "classical picture", essentially the family of properties accounts.

extension – a set or collection with strict membership – has been assumed, the jig is already up: To assume that a predicate has an extension in every possible world just is to assume that the meaning of the predicate is given by a property in the intensional sense. [20] The assumption of properties as predicate meanings has been assumed, if tacitly, at the outset.

How a counterextension might be split up isn't what matters. Supervaluationists might urge that a counterextension should be divided between the (super)false and the others. What matters is whether the extension and counter-extension are sets or like sets in having completely specific membership. If there are three sets, the extension and two further sets the union of which is the counter-extension, vagueness is gone just as much as if there are only two sets.[21, 22]

The villain of the piece is any properties account of predicates according to which the semantic value of a predicate either is or is something that determines an intension. In many, many places, Williamson talks of being thin as a property, and likewise being a heap, being bald, being mountainous…. A particularly explicit example: "On the epistemic view of vagueness, vague predicates stand for properties that each thing has or lacks." (1996, 333).[23]

Throughout his writings on vagueness Williamson appears simply to assume a properties account. In another passage Williamson attempts to argue for such a claim. At (1992, 147; see also 1994a, 196–197) He compares the possibility of bivalence failure in the case of vagueness with bivalence failure in the case of reference failure. When 'this dagger' has no referent, bivalence plausibly fails for 'This dagger is sharp.' But this is because "'This dagger is sharp' says nothing that could have been true or false…." (1992, 147).

Williamson dismisses any such way in which bivalence could fail in the case of borderline statements. He considers "a skeptical view" according to which

> [V]agueness is itself a kind of reference failure. Adjectives refer, if at all, to sharply defined properties, but [on this skeptical view] a vague one like 'thin' fails to single out such a property and so fails to refer; sentences of the form 'a is thin' say, strictly, nothing, whether or not a is a borderline case. (1992, 149; See also 1994a, 196–197)

Note how in this passage Williamson so clearly presupposes that "referring" to a property or being empty are the only alternatives.

Here is Williamson's response to this "skeptical view" (1992,149): "Since almost all our utterances involve vague terms, this view makes almost all of them mere noise." Consequently

[20] Section 3.4 will provide a more detailed discussion of how predicates can function without strict extensions.

[21] How then is the supervaluationst account an account of vagueness? The vagueness is smuggled in in the vague qualification that the precisifications be "appropriate" so that no robust understanding of vagueness has been provided. I will address this lacquna in Part II.

[22] Thus the perceived need for higher order vagueness. When we see in Part II how to understand this material with no appeal at all to sets, the perceived need for higher order vagueness will dissipate.

[23] See also the discussion below.

> To deny bivalence for vague sentences while continuing to use them is to adopt an unstable position…. Rapid alternation between perspectives inside and outside the practice [of treating statements as more than mere noise] can disguise, but not avoid, this hypocrisy….

In short, for utterances to function as more than mere noise the predicates used in these utterances must refer to "sharply defined properties." Once again we see that properties as predicate meanings has been simply assumed at the outset.

To summarize: Williamson's argument for bivalence assumes the Tarski biconditionals, in turn argued by appeal to his default principle, that in turn presumes that the meaning of predicates are "sharply defined properties": The properties are presumed from the beginning. But as everything turns on the presumption that the meaning of predicates are sharply defined properties, none of the intervening steps were necessary: From this presumption epistemicism is immediate. If the meaning of *P* is a sharply defined property then a predication, *Pa*, is true just in case *a* has that property, false otherwise: the presumption of sharply defined properties as meaning of predicates immediately gives bivalence. Then we get epistemicism by appending the interpretive gloss that, in this situation, the phenomenon of vagueness is not semantic and can only be epistemic, as discussed in the introduction.

Except for the question begging appeals to the default principle and the claim that otherwise much of speech would be "mere noise" Williamson offers no argument for a properties account. Many will say that no argument is required as the presumption of properties as predicate meanings is such a basic part of much thinking about semantics. What is really useful about Williamson's work is that it calls our attention to the overlooked circumstance that this "basic part of much thinking about semantics" has epistemicism as an immediate consequence. Either epistemicism must be embraced or this commonly assumed element of semantics has to be revised. In the next section I review and extend arguments showing why epistemicism has to be rejected. In part II I will propose a theoretical development that will guide such revision without falling into another trap than many see as equally problematic.

2.4 Difficulties with Epistemicism

Many have cited "the incredulous stare" and otherwise expressed astonishment at epistemicism.[24,25,26] In this section I will press the reasons for denying Williamson's repeated claim that use fixes properties, and through them completely precise extensions and counter-extensions. The difficulties leap to clearest relief when we press whether

[24] Again, leaving critical discussion of Williamson's margin for error principles to prior work of others.

[25] Machina and Deutsch (202, 27) briefly mention a few of the considerations in this section

[26] Caie (2012, 59; 2014); Schiffer (1999, 492, 493, 497); Keefe (2000, 64); Keefe and Smith (1999, 18, 21); Burgess (2001, 507); Machina and Deutsch (2002, 27,35); Ludwig and Ray (2002, 440–441; 2017, MS 14) Field (2000, 6); Wright (1995, 156), Tye (1997, 248–249), Magidor (Preprint, 9) Dorr (2003); Ebbs (2001); Enoch (2007); Gómez-Torrente (2002); Graff (2002); Horwich (1997); Kearns and Magidor (2012); López de Sa (2006); MacFarlane (2016); Magidor (2019); Sennet (2012). Papers listed in the references not otherwise mentioned in the text are included because they contain further valuable critical examination of Williamson'epistemicism.

a predicate's meaning fixes, for every actual and counterfactual context of use, a completely precise extension.

For meaning to determine the exact extension of 'bald' meaning will somehow have to take into account not only the number of hairs, but hair thickness, scalp-distribution, and doubtless much more. Similarly for 'heap', meaning-determining mechanisms must take into account size and shape of constituent parts and overall organization (or lack of it) of the collection: The number of constituents needed to make a heap will vary enormously among sand, pebbles, bricks, firewood…as will the arrangement. I can transform my disorganized heap of firewood into a well-organized stack. If it will be difficult for meaning-determining mechanisms to fix an exact extension for 'bald' and 'heap' relative to a context think how difficult this will be for terms such as 'funny' and 'suspicious'.

Let's look at an easier case, 'flat'. Following Unger (1971, passim.), if 'flat' is understood as perfectly, geometrically flat, nothing is flat. Williamson will not want to say that all our (positive) uses of 'flat' are false. How, then, is 'flat' to be understood? We often use 'flat' comparatively. Just how flat is that table? Is this table flatter than that one? What is involved when one says, simply, that something is flat is to be understood as "flat enough" for current concerns; specifically, that the differences between geometrically flat and the current case are negligible relative to our current interests and our standards for satisfying those interests. For meaning determining considerations to fix precise extensions for uses of 'flat', the meaning-determining mechanisms must provide a function from each individual or group's interests, standards and other contextual considerations, in all actual and counterfactual circumstances, to the exact extension of 'flat' for those contextual considerations. That's a lot to expect.[27]

There is a particularly virulent version of the kind of complaint in question that has been widely made in the literature. Williamson has offered a response. Let's look at both. [28]

How could meaning-determining mechanisms somehow fix the exact boundaries for things that are indicated in a very open-ended way, such as: "It's raining here" (Schiffer, 1997, 942), "Betty was standing roughly there", 'I worked for a little while yesterday' (Schiffer, 1999,488, 493); particles like: 'approximately', 'roughly', 'almost', 'not quite', and so on. (Wright, 1995, 153,154) In (1997a, p 953) Williamson attempts to address cases such as Jane's saying, "It's raining here":

> Schiffer objects that when one uses the demonstrative 'here' with ordinary vagueness, one has no way of identifying or picking out the sharply demarcated region to which one is referring, on the epistemic view…. It would be

[27] One commentator asked whether Williamson could appeal to reference magnetism (The idea goes back to Lewis, but see Sider, 2011). I do not here have space to defend my skepticism about this idea. In any case it is a notion that has been suggested for natural kind terms only and would fail completely for the problem we consider next.

[28] Strictly speaking, above I have only argued that epistemicism for predicates is the inevitable consequence of any properties account of the meaning of predicates. The problem cases to follow all involve reference. There is much to fill in here, but it will be clear enough that for the kinds of cases here in question, the issues will be parallel. In my (2018b, Sect. 3) I argue that the problems for reference are also ubiquitous.

> unreasonable to require… the speaker [to] know how to trace its boundaries in practice (consider 'this galaxy'). We should like an answer to the question 'Why does the demonstrative refer to x?'. Presumably, a central part of the answer will often be that the speaker is perceptually attending to x. But Schiffer gives no reason to suppose that one can perceptually attend to x only if one can locate the exact boundaries of x…. I can perceptually attend to the region with exact boundaries b even if no one is in a position to know that I am perceptually attending to the region with exact boundaries b; what I know is that I am looking at this region here.

This does not address Schiffer's point. What Schiffer wrote was that, broadly, for demonstrative reference to work, "I must… have some way of identifying my referent." (1997, 943) Williamson then responds by distinguishing between what one's perceptual system can attend to and what the agent can explicitly demarcate, claiming, with no argument, that the former can be a "region with exact boundaries b" even though the agent can't explicitly say just what those boundaries are. But Schiffer's point was, obviously, that if the referent of the use of 'here' is some completely precise area, the speaker *or* the speaker's perceptual system, *or something* going on in the act of demonstration has to identify that referent, has to pick out the area demonstrated from the uncountably many alternatives. We are owed an account of how that could possibly happen. In the passage under consideration, Williamson completely sidesteps that question.

In his (1999) Williamson again attempts to respond to Schiffer's version of the problem of how use of open-ended expressions could fix a place, time, or the like from the uncountably many alternatives. At (1999, 513) Williamson gives no more specific answer to this question than that, in cases like the ones under discussion, reference is fixed by speakers' intentions, going on to explain why this claim is not undermined by the fact that use can vary between speaker and hearer and that a hearer may take reference to be determined by deferring to the speaker's intentions. In other words, instead of responding to the challenge, Williamson changes the topic. In this passage Williamson also claims that the cases in question pose no special problem for epistemicism:

> But it is a mistake to suppose that epistemicism multiplies the candidates more than other theories of vagueness do. For example, if reference can somehow be indeterminate, many candidates will differ slightly from each other in their areas of indeterminacy. (1999, 513)

Williamson has interpreted 'reference is indeterminate' as reference is to something indeterminate. Then, if specific candidates for reference are indeterminate, many candidates will "differ slightly from each other in their areas of indeterminacy." Williamson concludes that any other coherent account of vague reference must take there to be as many distinct but "indeterminate" candidate referents as epistemicism's candidate determinate referents. In Part II we will see many ways in which "reference can be indeterminate" other than reference being (determinately) to something indeterminate.

I submit that, at least in the passages I have considered, Williamson completely fails to address the problem raised by open-ended use of demonstratives and particles such as 'approximately', 'roughly', 'almost',….

Next I discuss difficulties with Williamson's efforts to assuage worries about how completely precise extensions could be fixed.

The burning question: How could meaning be set so exquisitely so as to determine, in all actual and counterfactual borderline cases, exactly who counts as bald, what counts as funny…. Williamson's answer: Meaning is determined by use.[29] To track Williamson's response I divide the question into two parts: How does use determine meaning, and why should we think that the meaning thus determined is so precise that it invariably delimits completely precise extensions? Williamson responds to the second question by offering an answer to the first.

Here is how this conflation unfolds. Williamson has presupposed that "Words mean what they do because we use them as we do." (1994a, 205) The argument that I am examining is framed as a response to the objection that

> …if nature does not draw a line for us [as is plausible in the case of natural kinds], then a line is drawn only if we draw it ourselves, by our use. So (it is held) there is no line, for our use leaves not a line but a smear.' (1994a, 206).

Here is Williamson's response

> … 'drawing' is just a metaphor for 'determining'. To say that use determines meaning is just to say that meaning supervenes on use…. (1994a, 206)

What is the argument? The challenge was: How can use determine meanings *so precisely* as to determine all extensions exactly? The answer we are given is: To say that use determines meaning comes to saying that meaning supervenes on use, which is a response to the question, how does use determine meaning, but no response to the question, why should we think that the meaning so determined does better than "leaving a smear". Keefe (2000, 80 – 1) puts the problem this way: 'The fact that there can be no difference in meaning without a difference in use does not fix the boundaries of extensions any more than a pass-fail divide is fixed by the requirement that qualitatively identical exam scripts should receive the same mark.'[30]

[29] By 'use' Williamson understands dispositions for use of a term, not, or not just, the way the term has actually been used. (1994a, 205, 206, …) Williamson also specifies that the environment can be a "constitutive factor in meaning." He cites as an example the ways in which environmental facts about natural kinds can interact with use in determining meaning for natural kind terms. (1994a, 205-6) It would be natural to include environmental factors that fill in otherwise free parameters in statements (time, place, reference class…). Williamson gives no other specification of what kinds of environmental factors might be relevant. If just anything can count, the thesis is in danger of being trivialized. So I will assume that the only environmental factors to be included are ones that function in picking out natural kinds and the fixing of parameters in statements. Inclusion of a limited number of further environmental considerations will not affect what follows.

[30] Cf also Keefe and Smith (1999, 22), Keefe (1995, 394). Keefe (2000, 75–84) gives a general discussion of the problems with Williamson's views on use determining meaning and exact extensions.

Williamson continues his discussion of the changed question, how or in what way does meaning supervene on use:

> Although meaning may supervene on use, there is no algorithm for calculating the former from the latter' (1994a, 206, cf. also 209)

and

> The inability of the epistemic view of vagueness to provide a successful recipe [for calculating meaning from its supervenience base of use] is an inability it shares with all its rivals. Nor is there any reason to suppose that such a recipe must exist. (1994a 207)

But, just as before, the question was not, is there any recipe for getting from use to meaning, but is the meaning claimed to supervene on use one that will fix extensions exactly?

So far we have seen that no reason has been offered for thinking that the meaning determined by use should have the claimed precision. A further worry is: What is the mechanism through which use determines one meaning as opposed to some other? The response was that, here, 'determined by' is to be understood as "supervene on". The new worry is that the "determined" of "supervenience" is the wrong kind of determination. A subvening domain determines the supervening domain only in the sense that there is no variation in the latter without some variation in the former. But that's consistent with there being no interesting sense in which the subvening "makes" the supervening be what it is as opposed to something else. If use determines meaning, we expect that there be some way, something about the situation, that results in the postulated exact extensions coming out one way rather than another. Appeal to no more than supervenience leaves this all a complete mystery.

In a way Williamson concedes the last complaint. He writes that 'Meaning may supervene on use in an unsurveyably chaotic way. (1994a, 209) But if the basis is unsurveyably chaotic, it is a mystery how people could learn meanings or reliably use language in communication.

Another twist in this tangle: One can agree that (in the weak duplication sense) meaning supervenes on use. But when it comes to how things turn out, meaning guides use at least as much as use guides meaning. We use words the way we do in large part because of the meaning we take them to have. In the sense of 'determine' of interest here, if anything, meaning determines use more than use determines meaning.

Turning to yet one more general problem with Williamson's account: For a given term and time, which uses – that is, what dispositions to use a term – are relevant in fixing its meaning? At (1994a, 211) Williamson puts the issue in terms of communities of language users. Since, for an unclear case, we don't know whether the case is in the extension that is set by the meaning of a term, it might seem that epistemicism "prevents us from knowing what we mean." (209) Williamson's response: "On the epistemic view, our understanding of vague terms is not partial. The measure of full understanding is…complete induction into a practice." (1994a, 211) Leaving aside perplexities about how to understand "complete induction into

a practice," clearly Williamson intends (1994a, 211) that such a practice is the practice of a linguistic community. That is, it will be the use in such a community that sets the meaning that, in turn, sets the completely exact extensions of vague terms.

Not only is it left free-floating just which vocalizations count as meaning-fixing uses,[31] just what community is relevant?[32] Williamson gives us nothing to guide us as to how to approach this issue. At (1994a, 211 – 12) Williamson explicitly declines to rule out cases where one individual can count as the relevant community. For larger communities it will often be arbitrary exactly who gets included. There is no specification of which communities are relevant, and so no specification of which uses fix the extensions. This is particularly problematic since, in (1994a, Ch 8) Williamson so clearly commits himself to the claim that very small changes in use will result in small changes in meaning and so, usually, in the relevant extensions. To be clear, the complaint is not that Williamson has not given a detailed account of what the communities are. The problem is that we have no idea how to start on such a project, what could exactly fix the communities. The communities are, at least in details of membership, arbitrary.

In (1997a) Williamson appears to try to address this worry. At (1997a, 952) he acknowledges that "The notion of a speech community is itself vague." Consequently "[a] sorites paradox threatens." There follows a long paragraph that supports the conclusion that

> The sorites paradox does not arise, because in a given context not all slight differences in use imply sameness [did he here mean differences?] in speech community: the implication holds only for some slight differences that are salient in that context.

Even if we change the 'sameness' to what must have been intended, 'differences', it's hard to see how this conclusion addresses the issue of who counts as members of the relevant use-and -extension-determining linguistic community.

Not only do we have no idea how use could possibly fix the exact extensions from uncountably many alternatives and in innumerable different possible contexts. In addition just which uses are relevant for a given term is left open-ended.

3 Part II: Prolegomena to an Alternative to Properties Accounts

In Section 2.3 I argued that Williamson is presupposing a properties account of predicates, which makes his view inevitable. Most of Williamson's discussions are efforts to accommodate the insuperable consequences. I have extended the many extant examinations to demonstrate that these efforts fail completely. What I take epistemicism's difficulties to show is that epistemicism's underlying assumption must be replaced. What is involved in this replacement, and why, will become clear

[31] The only response that I have found to this worry is that the relevant uses are the ones that are "are salient in that context." (1997a, 952) See just below.

[32] Morton (1995 p. 275) briefly makes a similar point.

by placing the assumption in its frequently used setting, compositional semantics. We will see why taking predicates to express some kind of properties has seemed so natural, why alternatives have been rejected by many, and how, despite being a disaster when it comes to understanding vagueness, a properties account of predicates nonetheless has a useful role to play in semantic theory. The role of a properties account in this setting will, in turn, point to a range of candidate alternatives. Many of these candidates are well known and widely discussed, and also by many dismissed out of hand. The analysis will illuminate such criticism and make clear why, in important respects, dismissal has been over hasty.

3.1 Theory of Attachment of Language to the World and Theory of Semantic Structure

The point of departure is the commonly used notion of compositional semantics. The meaning of expressions are build up following the structure of the syntax of a sentence. One starts with root meanings or "semantic values" for root syntactic units, by which I mean the units that get their semantic interpretation through our methods for applying language to the world.[33] These root values then combine following syntactic structure to form more complex semantic values, issuing in the semantic values for whole sentences, in the case of declarative sentences often called (structured) propositions. The present point focuses on the contrast between the compositional aspect and attachment to the world. The smallest, root values have to be attached to the world – treated by what I will call a "theory of attachment". The compositional structure then combines the root semantic values – treated by what I will call a "theory of semantic structure".

I want to emphasize that two distinguishable theoretical topics are in question. We can in many ways separate them, and there is much to be gained by doing so. To separate the theoretical problems we put in idealized placeholders for the root semantic values. Treating the root semantic values as referents and properties then functions as a way of bracketing the problem of ultimate attachment to the world. Talk of referents and properties work as placeholders for the complex ways in which ultimate attachment to the world takes place, to be addressed on another day or by another discipline. Details of how attachment occurs are generally not relevant to compositional structure.

It is easy mistakenly to take these idealized placeholders literally. Instead of treating them as idealized placeholders, often, and often in ways not explicitly expressed, the root semantic values for referring terms and predicates are taken to be "things in the world",[34] referents and properties (again, always understood to include relations). There is a unique thing, the Eifel tower, that gets attached as the referent of 'Eifel tower'. It is often also

[33] These smallest units can be smaller than lexical entries that themselves can have structure with semantic content: gender, case, mass vs. count nouns and many other features for which lexical entries are marked that have both semantic content and syntactic effects.

[34] E.g., see Fodor and Lepore (1995, p. 255).

assumed, if tacitly, that there are precise properties, such as the color red and being mountainous (Williamson, 1994a, 268) that get attached to the terms 'red' and 'mountainous'. (I use these two as examples because, when one stops to think about it, these obviously do not attach to specific properties under any reasonable reading of 'properties'.)

The methodology to which I am appealing, separating problems by putting in idealized "placeholders", is a commonplace in the physical sciences but may be less familiar in other disciplines. One commentator insisted that nothing short of a fully integrated account is acceptable. This would be like rejecting a physicist's use of the simple law of friction unless accompanied will a full account of how frictional forces arise at the atomic scale. Not only is a fully unified account generally not available, depending on objectives the details can obscure more than they illuminate. Another commentator objected that it is "extremely hard to see how precisely [the proposed methodology] generates a compositional account of meaning." It doesn't. Rather we start with current theories of compositional semantics that appeal to something called "properties" as root semantic values. We then withhold any literal reading, instead taking talk of properties as an idealization that puts to one side the difficulties in understanding the attachment of root semantic values to the world and understand the latter as a separate problem.

The problem of attachment has been addressed in a variety of ways that are often characterized as "use theories". Problems arise when use theorists, seeing that their ideas look promising for the point of attachment to the world jump to the conclusion that this kind of theorizing will suffice for a full understanding of language, including the parts that have to be understood in terms of compositional considerations. In opposition, opponents of use accounts, seeing that use approaches won't suffice for a complete understanding, conclude that the ideas have no application.[35] Both attitudes are mistaken. The problem of worldly attachment and the problem of composition are very different sorts of problems, usefully separated by the methodological device I've described. Once separated the two problems need treatment with entirely different accounts.

3.2 The Approach of Ludwig and Ray

Ludwig and Ray (to appear) provide a prior proposal that agrees in many respects with the view that I am developing. According to Ludwig and Ray terms in a language come with "rules of application":

> A predicate is semantically complete only if its [complete] rules of application (positively) determine whether it is true of, [or] false of,….anything or n-tuple to which it is applied. (Preprint, 6–7)[36]

[35] Fodor and LePore (1995) provide a good example of this mistake. Kamp and Partee (1995) illustrate the kinds of contortions that occur if one insists on treating attachment and composition together.

[36] At the ellipsis I have omitted their third condition, "or neither true nor false of". Ludwig and Ray have explained to me (personal communication) that in their exposition this alternative needs to be included because they want the rules of application of a semantically complete predicate to cover cases of bivalence failure when sortal requirements are not satisfied, an entirely different source of bivalence failure than vagueness. I omit this condition because, following Williamson, I am putting to one side complications of sortals as they do not importantly bear on the issues pertaining to this discussion.

A predicate has an extension and counterextension iff it has "complete rules of application". (Preprint, 9–10)

Ludwig and Ray then press sorites arguments into an argument for the conclusion that

> …'bald', and similarly for any vague expression, does not have an extension and counterextension because it does not have complete rules of application, that in turn comes to it not having a 'complete meaning' (Preprint, 9, 10)

Saying that 'bald', and with it vague predicates generally, have no extension, at first blush seems absurd. But keep your eye fixed on the fact that here extension – what vague expressions do not have - is understood in the technical sense from compositional semantics. Extensions are theoretically required to be sets or collections with well determined membership, not extensions in any informal, pretheoretical sense. In the next section I will discuss how predicates can function without extensions.

What, then, do Ludwig and Ray make of vague terms having no extensions?

> [A] meaning introducing intention [may] not specify a plan for its use that suffices to fix an extension, and, hence, it does not suffice to fix a meaning (or a complete meaning, if you like). Still, it can induce a practice in the use of the expression. And we can, if we like, pretend that, or act as if, the practice has been filled in, because we have an idea about where safe areas of operation are because we know at least roughly what [are] the intentions for the term left unspecified. (Preprint, 10)

So we

> Ordinarily operate under the pretense that our terms are semantically complete, and since we are not unaware of where our practices give out, this usually gives us no trouble. (Preprint, 12)

Communication is then

> highly pragmatic, but this seems, in fact, exactly right. [Things go smoothly when] a use is well within the standard practice, that is in turn characterized in part by when it is not safe to use terms in application to objects or a range of objects because the practice doesn't give us much guidance in how to use them. (Preprint, 15)

I will restate what I take to be the same or very similar ideas in the terms that I have introduced and will use in the remainder of this paper. We use terms with inexact (Ludwig and Ray – incomplete) methods[37] of application. Learning to use these methods involves learning where they apply safely, not so safely, and very badly. When we reliably discern that the context allows safe application we harmlessly engage in the idealization (Ludwig and Ray – pretense) that the methods are precise,

[37] Ludwig and Ray use the terms 'plan' and 'practice'. It is not assumed, indeed it rarely happens, that these are applied self-consciously. Rather what is in question here are dispositions for term application that theorists can characterize in terms of regularities or open ended-rules.

that they work by specifying specific referents and properties as treated in our highly idealized theory of semantic structure.[38]

One upshot is that statements using vague terms are *never* either true or false, that is, not true or false in the idealized sense that is in question in semantic theory, and tacitly in much of philosophy. In this idealized sense of truth, predicating a property of a referent is true just in case the referent has the property. But, cases in mathematics aside, these properties are, one and all, idealized creatures.[39] Instead, treating the phenomena of vagueness requires recognizing that these idealizations of semantic theory are the wrong theoretical tools for that job. The vagueness of terms is the linguistic face of the inexactness of the methods that govern their inexact application. Consequently, the right places to look to understand vagueness are the incomplete or inexact methods that govern the application of vague terms.

In his discussion of nihilism (1994a, 165 ff.) Williamson rejects "usefulness" and "pretense" accounts. Williamson defines nihilism as the doctrine that all vague terms are meaningless. (1994a, 165) and then argues that it is absurd for nihilists to substitute 'useful' for 'true' (1994a, 169) because 'common sense belief' and 'useful' are vague, and so on the nihilists' own account the statements they want to make are not just not true, they aren't even meaningful. (1994a, 169) Williamson likewise dismisses any suggestion that

> '[w]e are to pretend that our words have precise meanings although there are no precise meanings that we are to pretend that they have…. But for the nihilist the words with which such a kind might be delineated are empty.' (1994a, 170–171).

On Williamson's definition of nihilism, since nihilists take all vague terms to be meaningless nihilists literally cannot make such assertions.

Williamson's definition of nihilism isn't arbitrary. Given Williamson's larger position this definition is natural, really is forced. At the end of Section 2.3 we learned that, when one chases down Williamson's grounds for claims, and the grounds for grounds…, the trail ends with the assumption that "Adjectives refer, if at all, to sharply defined properties." So statements made with a predicate that does not refer to a sharply defined property "say, strictly nothing" (1992, 149) That is, according to Williamson the only alternatives for a predicate are to have a property as its meaning or to be meaningless – "empty".

In effect this claim was supported by appeal to the Tarski biconditionals, but the justification for them then turned out to beg the question of whether having properties as meanings is the only way a predicate can be meaningful. Taking truth to be defined by the biconditionals, the project of looking for other ways in which

[38] It is extremely misleading to label such idealized use, or pretense as "fictional". That a representation is not completely accurate does not automatically make it a fiction! In Winsburg's example (2008, 180), if a biography is found to have some mistakes in it, we do not then put it on the fiction shelves of a book store. Winsberg explains that what determines whether inaccurate representations are appropriately classified as fictional or faithful, though imperfectly so, turns on their intended use. (2008, 180-1)

[39] In my (2017) I have argued that there are no properties attached to 'red', 'funny'…. In my (2018a) I show that even attachment to physical properties and quantities fails in physics.

predicates can function meaningfully requires that we give up truth in its classical and exacting sense. Either we drop evaluation in terms of truth and falsity altogether in favor of other standards such as ones having the form: Have the inexact and incomplete methods of application been followed well enough to ensure reliable communication and to satisfy community standards? Such standards are inexact and open ended. But it no more follows that they are not real standards than it follows that laws governing our behavior are not real laws because they often require judges to interpret them. Or we fashion a more open-ended and flexible notion of truth that functions similarly to the forgoing and that more faithfully characterizes 'true' as we use it in everyday life.[40]

Once one refrains from begging the question and allows other kinds of ways in which a term could be meaningful, it is easy to explain the kind of precise meaning that is pretended. The pretense is that, counter to fact, the word in question stands for a sharp property which is exactly the way I have suggested that we think of root semantic values where I used the term 'idealization' rather than 'pretense'.

3.3 Use Accounts and Truth

On properties accounts of predicates a simple predication, *Pa*, is true if *a* has the property expressed by *P*, false if *a* does not. On use accounts simple predications are neither true nor false, at least not in this traditional, and in my view idealized, way of thinking about truth.[41]

How then, are we to make sense of ordinary attributions of truth and falsity to ordinary claims using vague terms? The larger position needs a different way of thinking about truth that I have developed in a different publication (Teller, 2017), the results of which I will summarize here. The classical, correspondence notion of truth, that for this paragraph I will write as 'Truth', is retained, but only as an idealization, corresponding to the idealization involved in taking the semantic value of a predicate to be a property: *Pa* is True just in case the referent of *a* has the property taken to be expressed by *P*. A second notion of truth is introduced, that for this paragraph I will write as 'truth'. A predication counts as true just in case it functions as a Truth for any purpose that it is reasonable to think might come up. Truth and truth are closely connected. On the one hand I appeal to Truth in explaining "functions as a Truth." On the other hand Truth is the idealized, limiting case of truth. Note that Truth is a precise notion, truth a vague one. To see the need to take truth to be vague, compare

(1) 'Harry is bald' is True iff Harry is bald
(1') 'Harry is bald' is true iff Harry is bald

where Harry is a borderline case of being bald. On any account of vagueness that gives borderline cases any sort of truth-value other than True or False, (1) is False,

[40] See Teller (2017) and Elgin (2017).

[41] This is also the conclusion of Braun and Sider (2007), Ludwig and Ray (2002, to appear) and Teller (2017).

or perhaps defective in some other way. (1') will fare better if we take 'truth' to be vague in a way that reflects the way vagueness is treated in borderline cases.

These considerations clear away possible worries about logic. When circumstances are sufficiently clear cut to allow simple predications to be treated as having determinate truth values, use of classical logic will not mislead. For borderline cases one is ill advised to apply classical logic anyway. This is just as in Section 2.2's response to Williamson's worries about logic.

3.4 Meaningful Predicates Without Extensions

Ludwig and Ray say that predicates with no "complete meaning" don't have extensions or counter extensions. How could that be? I will address this in more detail in another paper. But here is the general idea:

Following Frege, the idea of the meaning of a term is of that which determines to what the term applies. Conventionally we think of such as some objective property, something independent of us, something "out there in the world." And, on pain of some kind of ontological vagueness, such will have a completely determinate extension and counter-extension.

Taking predicate meanings as objective extension-determining properties is not yet a full account of speakers' use of these predicates. Speakers have to make judgements as to what a predicate should be taken to apply. On any view according to which it is some kind of property that intervenes to determine whether something goes into an extension, speakers would, at least in effect, be making judgements about whether or not a referent has the property in question. In judging whether 'red' applies to a barn a speaker would, on such conventional accounts, be, at least in effect, judging whether the barn has the property of being red.

On conventional views, such an intervening property is functioning as a middleman between a speaker's method of application and the things to which a predicate is applied. It has been the burden of Part I that postulating such middlemen leads to the absurdity of epistemicism. Any such middlemen must be rejected. In other places (Teller, 2018a, b) I have independently argued that attachment to such middlemen can never occur simply because the world is too complicated. As I have urged above, thinking in terms of these intervening properties is a simplification, an idealization, often extremely useful when not mistakenly taken literally.

We are left then with the methods that speakers use to apply predicates, Ludwig and Ray's "practices". Meaning holists will take all these methods to count as contributing to predicate meaning. There are also various options for taking some of these methods to contribute to predicate meaning while other methods have been contingently found to track the meaning-conferring methods. This is not the place to sort through the many options. What matters here is that all the methods, in particular the meaning conferring ones, must work in a world too complex for any finitely stateable or humanly learnable rule to operate decisively for absolutely every case. In marginal cases speakers must use their discretion, which will be influenced by current interests and all manner of varying contextual considerations. So, for marginal cases, speakers

may faultlessly disagree. As long as such disagreements occur infrequently, language still operates smoothly.

Consequently, these meaning conferring, open-ended rules of application don't determine extensions, sets or set-like in having completely determinate membership. The open-ended rules of application for a predicate leave some cases for which there is no speaker independent fact of the matter whether the predicate correctly applies or not. Nor can there be a set of cases for which the rules of application *always* yield the same determination, that is, never allow any speaker discretion. No doubt there are many such cases but no set of such cases. How could there be such cases but no set of them? Any elusive boundary between the cases in which a rule determines application and when it doesn't, speaking metaphorically, "moves around" with changing circumstances of application resulting in, or resulting from, varying interpretations of the rule. This metaphor can be cashed in by observing that any envelope of cases for which there can be faultless disagreement would have to cover all possible counterfactual cases, but this collection of cases isn't well characterized. When the counterfactual cases become more and more extreme it will no longer be clear whether we are talking about the same word with the same meaning.[42]

I have argued that the theory of semantic structure, addressed to very different issues than a theory of attachment, will treat the root semantic values as idealized referents and properties. Vagueness cannot be understood within the idealizations of conventional semantic theory. Instead vagueness has to be understood by studying the inexactness of the rules or methods with which we apply terms to the world. I will not propose any one specific theory of attachment or add detail to the many accounts that already have a foothold. Rather I will canvass some options with a few comments on how they might provide insight.

A caveat: We must not expect one uniquely correct account! Most broadly, theories of attachment work out how we classify things. Things in different domains may need very different methods of classification, as will become clear with the examples below. So we should expect that these complex phenomena will be addressed by a range of accounts that complement one another. Different domains of application may proceed in quite different ways. Each account is to be evaluated for its strengths and weaknesses, keeping only those that do well for a robust range of phenomena.

The most obvious candidates are prototype and exemplar theories, accounts that work by applying a similarity metric to prototypes or exemplars, to be include here also cluster concept accounts that work by applying a cluster of concepts that may be weighted and where the language user is allowed some leeway in what components to include and how to weigh them.

We often classify things by the roles they play in our lives. These might be alternatives to prototypical characterizations, or they might be combined with some prototype account as a guide to what characteristics should go into the relevant prototype. Things are often classified functionally: knife, chair, sled, nurse… Classification may be by the

[42] In many cases it will be easy to specify strict semantic relations between meanings given as open-ended rules of application. For example, with converse relations, such as *south of* and *north of*, there would be two specifications of the same open-ended rule of application with switched variables.

function an object serves but also by the way that function is achieved. For example consider the difference between a sled and a toboggan. There is a nice lesson from the case of 'nurse'. Twenty years ago and before, the example of 'nurse' appeared in the literature critical of prototype approaches. At that time we expected nurses to be women. Consequently, it was suggested, on a prototype account the expression 'male nurse' seemed problematic. The problem was that analysts conflated the basis for classification – I suggest in such cases it should be function – with practical means of identifying instances. Often we can make a good distinction between a basis for classification – a characterization or "definition" – and epistemic means of identifying instances.

Insofar as a term heads a category that is characterized functionally, the term will unavoidably be vague. One ineliminable source of indefiniteness, and so vagueness, is that the function has to be played "sufficiently well", which is amorphously context dependent, and ultimately up to speaker and hearer to decide whether the role is played sufficiently well to support communication and other interests in that context. This idea will reoccur in the further examples.

Things are also classified in terms of the sensations they cause: colors, sounds, tastes textures…. There is no need or basis for supposing that there must be uniquely correct conditions for application of such terms. It suffices to suppose that speakers and hearers have guidelines - Ludwig and Ray's "practices" – but use their contextually guided discretion where to draw lines. As above the test will be whether the needs of communication and other interests are sufficiently satisfied.

We classify things in terms of the emotions or other mental attitudes that such things tend to bring about: funny, exciting, frightening. 'Beautiful' is an instructive example. Ultimately its use is guided by the tendency to elicit one or another aesthetic response, but in ways that are highly variable. The actors must judge whether the relevant responses are appropriate for the people and context that is in question.

All of these examples can be seen as falling under so called "use" accounts. Expositions of use accounts are often distressingly unspecific – sometimes even less specific than my brief remarks just above! But all these methods, or practices, or rules, or guidelines are susceptible to empirical investigation. An often neglected consideration that should be born in mind: Since such methods are always in support of communication and other interests, open-ended rules will generally allow speaker discretion in problematic cases. In particular, open-ended rules will generally allow speakers and hearers to slot in their locally applicable standards and preferences. The variability of these standards and preferences are a major source of the open-endedness – and so vagueness – of methods of attachment of language to the world.

3.5 Higher-Order Vagueness

The foregoing considerations now apply to show that higher order vagueness is, in almost every respect, an artifact of thinking in terms of predicates with exact extensions. The conclusions about higher order vagueness require their own detailed development. I can here only suggest the general approach.

Why does higher order vagueness seem forced on us? For example suppose that 'short' had no second order vagueness. Then there would be only (clearly) short people,

(clearly) not short people, and the unclear cases. But then we can say that all people are either short, or unclear, or not short. Whether or not we revise 'not short' to include those previously classified as unclear cases, vagueness is gone. Since 'short' *is* vague, there must be some "fuzziness" between the unclear cases and the others, cases for which it is not clear whether or not they are clear – the cases of second order vagueness. The pattern of argument now reapplies to generate all higher orders of vagueness.[43]

This argument is fallacious. It assumes that if there were no second order vagueness, there would be a set – a determinate collection – of instances that count as unclear cases. This is right if predicates express properties but need not be right on use accounts. On use accounts just who can be correctly classified as short "moves around". In "close calls" whether or not someone gets classified as short will depend on fine details of the circumstances, in sufficiently close cases, a choice may be arbitrary. All this in agreement with the open-ended methods governing the application of 'short' that accord discretion to language users in difficult cases. It is the contextual open-endedness of a term such as 'short' that does the work mistakenly attributed to higher order vagueness.

One might be tempted by a counter argument according to which, for a fixed reference class, there will at least be a greatest height so that no matter what the counterfactual circumstances, any one with that height or less would, in those circumstances, count as short. This counter argument makes the same kind of mistake, this time assuming clear delineation of what counterfactual circumstances are in question. But when the counterfactual circumstances become too extreme, the whole function of the methods for applying 'short' begin to break down, and it is no longer clear whether we are considering the same word with the same meaning. So, just as with the original argument for higher order vagueness, this rescue presupposes a determinate set of cases where there is none.

3.6 Conclusion and Moral for this Tale

A two sentence summary of the thesis of this paper: The intractability of vagueness is an artifact of all or nothing thinking. In particular, that the world is to be understood in terms of properties, some kind of characteristics that are completely sharp or, more broadly, by treating idealizations as if they were exactly correct accounts. Correcting this mistake does not thereby answer detailed questions about vagueness, but transforms them into real, broadly tractable empirical questions about application of language to the world.

There is also a moral to the tale: Thinking that there is always one right answer to any question is an endemic characteristic of western intellectual thought. Such thinking deflects us from appreciating that, often, there is a range of worthy approaches, different ones working well for different kinds of problems or cases.

[43] This argument is the same as, or at least very similar to, one given by Sainsbury (1991, 167–170 and 1999, 253-5). Sainsbury then points out that when we collect all finite orders of higher order vaguegueness for a term, we are right back where we started.- and can march on into the transfinite!

Declarations

Conflict of Interest The author declares that they have no conflict of interest.

References

Braun, D., & Sider, T. (2007). Vague, so untrue. *Nous, 41*(2), 133–156.
Burgess, J. (2001). Vagueness, epistemicism, and response-dependence. *Australasian Journal of Philosophy, 79*, 507–524.
Caie, M. (2012). Vagueness and semantic indiscriminability. *Philosophical Studies, 160*, 365–377.
Caie, M. (2014). Metasemantics and metaphysical indeterminacy. In A. Burgess & B. Sherman (Eds.), *Metasemantics: New Essays on the Foundations of Meaning* (pp. 55–96). Oxford: Oxford University Press.
Dorr, C. (2003). Vagueness without ignorance. *Philosophical Perspectives, 17*, Language and Philosophical Linguistics: 83-113.
Ebbs, G. (2001). Vagueness, sharp boundaries and supervenience conditions. *Synthese, 127*(3), 303–323.
Eklund, M. (2011). Recent work on vagueness. *Analysis, 71*(2), 352–363.
Elgin, C. (2017). *True enough*. MIT Press.
Enoch, D. (2007). Epistemicism and nihilism about vagueness: What's the difference? *Philosophical Studies, 133*(2), 285–311.
Field, H. (2000). Indeterminacy, degree of belief, and excluded Middle. *Noûs, 34*(1), 1–30.
Fine, K. (1975). Vagueness, truth and logic. *Synthese, 54*, 235–259.
Fodor, J., & LePore, E. (1995). The red herring and the pet fish why concepts still can't be prototypes. *Cognition, 58*, 253–270.
Gómez-Torrente, M. (2002). Vagueness and margin for error principles. *Philosophy and Phenomenological Research, 64*(1), 107–125.
Graff, D. (2002). An anti-epistemicist consequence of margin for error semantics for knowledge. *Philosophy and Phenomenological Research, 64*(1), 127–142.
Horwich, P. (1997). The nature of vagueness. *Philosophy and Phenomenological Research, 57*(4), 929–235.
Kamp, H., & Partee, B. (1995). Prototype theory and compositionality. *Cognition, 57*, 129–191.
Kearns, S., & Magidor, O. (2012). Semantic sovereignty. *Philosophy and Phenomenological Research, 58*, 322–350.
Keefe, R. (1995). Review of *vagueness* by Timothy Williamson. *The Philosophical Quarterly, 45*(180), 392–395.
Keefe, R. (2000). *Theories of vagueness*. Cambridge University Press.
Keefe, R., & Smith, P. (1999). *Vagueness: A reader*. MIT Press.
López de Sa, D. (2006). Is everything precise? *Dialectica, 60*(4), 397–409.
Ludwig, K., & Ray, G. (2002). Vagueness and the Sorites Paradox. In J. Tomberlin (Ed.), *Philosophical Perspectives: Language and Mind.* Blackwell.
Ludwig, K., Ray, G. (2017). Preprint. The Truth About Vagueness. Forthcoming in *On the Sorites Paradox*, eds. Ali Abasenezhad and Otavio Beuno, Springer, 2017.
MacFarlane, J. (2016). Vagueness as indecision. *Aristotelain Society Supplementalry Volume*, *XC*, 256–283

Machina, K., & Deutsch, H. (2002). Vagueness, ignorance, and margins for error. *Acta Analytica, 17*, 19–45.
Magidor, O. (2019). The epistemicist solution to the Sorites Paradox. In E. Zaardini, & S. Omes (Eds.), *The Sorites Paradox*, Cambridge University Press (CUP says on line version will be available August 2019).
Mahtani, A. (2004). The instability of Vague Terms. *The Philosohical Quarterly, 54*, 570–576.
Mahtani, A. (2008). Williamson on inexact knowledge. *Philosophical Studies, 139*, 171–180.
McGee, V., & McLaughlin, B. (1998). Review of *vagueness* by Timothy Williamson. *Lingistics and Philosophy, 21*(2), 221–235.
Morton, A. (1995). Review of vagueness by Timothy Williamson. *Philosophical Books, 36*(4), 272–276.
Ray, G. (2004). Williamson's Master Argument on vagueness. *Synthese, 138*(2), 175–206.
Richard, M. (2000). On an argument of Williamson's. *Analysis, 60*(2), 213–217.
Sainsbury, M. (1991). Is There Higher-Order Vagueness? *Philosophical Quarterly, 41*, 167–82.
Sainsbury, R. M. (1999). Concepts without boundaries. In R. Keefe & P. Smith (Eds.), *Vagueness: A Reader* (pp. 251–264). MIT Press.
Schiffer, S. (1999). The epistemic theory of vagueness. In *Philosophical Perspectives Vol. 13, Epistemology*, 481–503.
Schiffer, S. (1997). Williamson on our ignorance in borderline cases. *Philosophy and Phenomenological Research, 57*(4), 937–943.
Sennet, A. (2012). Semantic plasticity and epistemicism. *Philosophical Studies, 161*(2), 273–285.
Sider, T. (2011). *Writing the Book of the World*. Oxford University Press.
Teller, P. (2018a). Measurement accuracy realism. In I. F. Peschard, B. C. l. van Fraassen, (Eds.), *The Experimental Side of Modeling* Minnesota Studies in the Philosophy of Science, 21. Minnesota University Press.
Teller, P. (2018b). Making worlds with symbols. *Synthese, 198*(Suppl 21), 5015–5036.
Teller, P. (2017). Modeling truth. *Philosophia, 45*(1), 143–161.
Torrago, L. (1998). Vagueness. *Philosophical Review, 107*(4), 637.
Tye, M. (1997). On the epistemic theory of vagueness. *Philosophical Issues, 8*, 247–253.
Unger, P. (1971). A defense of skepticism. *The Philosophical Review, 80*(2), 198–219.
Williamson, T. (1992). Vagueness and ignorance. *Proceedings of the Aristotelian Society Supplementary Volumes, 66*, 145–177.
Williamson, T. (1994a). *Vagueness*. Routledge.
Williamson, T. (1994b). Definiteness and knowability. *The Southern Journal of Philosophy, 33*(Supplement), 171–191.
Williamson, T. (1996). What it takes to make a heap. *Erkenntnis, 44*(3), 327–339.
Williamson, T. (1997a). Reply to commentators. *Philosophy and Phenomenological Research, 57*(4), 945–953.
Williamson, T. (1997b). Imagination, stipulation and vagueness. *Philosophical Issues 8, Truth*: 215–228.
Williamson, T. (1999). Schiffer on the epistemic theory of vagueness. *Philsophical Perspectives, 13*, 505–517. Epistemology.
Winsberg, E. (2008). A function for fictions: Expanding the scope of science. In M. Suarez (Ed.), *Fictions in Science: Philosophical Essays on Modeling and Idealization* (pp. 197–189). Routledge
Wright, C. (1995). The epistemic conception of vagueness. *The Southern Journal of Philosophy, 333*(Supplement), 133–159.