

Epistemicism, Distribution, and the Argument from Vagueness

Ofra Magidor

(penultimate draft of paper *forthcoming in Nous*)

Abstract: This paper consists of two parts. The first concerns the logic of vagueness. The second concerns a prominent debate in metaphysics. One of the most widely accepted principles governing the ‘definitely’ operator is the principle of Distribution: if ‘p’ and ‘if p then q’ are both definite, then so is ‘q’. I argue however, that epistemicists about vagueness (at least those who take a broadly Williamsonian line) should reject this principle. The discussion also helps to shed light on the elusive question of what, on this framework, it takes for a sentence to be borderline or definite. In the second part of the paper, I apply this result to a prominent debate in metaphysics. One of the most influential arguments in favour of Universalism about composition is the Lewis-Sider argument from vagueness. An interesting question, however, is whether epistemicists have any particular reasons to resist the argument. I show that there is no *obvious* reason why epistemicists should resist the argument but there is a *non-obvious* one: the rejection of Distribution argued for in the first part of the paper provides epistemicists with a unique way of resisting the argument from vagueness.

Introduction

This paper consists of two parts: the first directly concerns the logic of vagueness, and the second applies to the result of the first part to a prominent debate in metaphysics: the Sider-Lewis argument from vagueness.

Different theories of vagueness disagree about many features governing the ‘definitely’ operator. However, one principle which is widely accepted by theorists of different camps (including epistemicists) is the principle of Distribution: if ‘p’ and ‘if p then q’ are both definite, then so is ‘q’.

In the first part of the paper I argue that epistemicists (in particular those who follow broadly Williamsonian lines) should reject Distribution. My discussion also sheds light on the notoriously elusive question of what, within this framework, is required for a sentence to be borderline.

This rejection of Distribution has various important implications: technically, it entails that the logic of ‘definiteness’ cannot be represented as a normal modal logic, and consequently, cannot be given a standard Kripkean possible-worlds semantics. More informally, it means that definiteness not only fails to be closed under entailment it is not even closed under *definite* entailment. In the second part of this paper, I demonstrate that the rejection of Distribution also has important implications to a prominent debate in metaphysics.

Universalism about composition is the claim that every class of objects has a fusion. One of the most influential arguments for universalism is the Lewis-Sider argument from vagueness. An interesting question is whether epistemicists about vagueness have any particular resources to resist the

argument. It is sometimes suggested that epistemicists can easily resist the argument from vagueness. I argue this impression is misguided (there is no obvious reason why a commitment to epistemicism should lead one to reject the argument). However, I show that there is a *non-obvious* reason for epistemicists to resist the argument: the rejection of Distribution defended in the first part of the paper undercuts the support for one of the premises in the argument.

The discussion in the second part also serves to reinforce the argument in the first part, by providing further potential counterexamples to Distribution.

Part A: Epistemicism, Semantic Plasticity, and Distribution

A1. Vagueness and Distribution

One distinction which theories of vagueness are concerned with is the distinction between *vague* and *precise* statements: while ‘ $2+2=4$ ’ is precise, ‘Harry is tall’ is vague. But another, perhaps more interesting distinction is that between sentences that are *borderline* and those which are not. Thus for example, if Harry’s height is 1.75m then arguably ‘Harry is tall’ is borderline, and thus neither the sentence nor its negation are definite. By contrast, if Harry’s height is 2m, then ‘Harry is tall’ is definite, not borderline (despite being vague).¹

Different theories of vagueness disagree about which principles govern the notions of definiteness and borderlineness. For example, according to many theories $B(s)$ entails that s is neither true nor false, while according to others, sentences can be borderline at the same time as being true or false.

Despite these disagreements, there are also some principles governing these notions which are widely accepted by competing theories. One such principle concerns the connection between borderlineness and definiteness: it is widely (probably universally) agreed that $B('p')$ iff $\neg D('p') \wedge \neg D('\neg p')$ and that

¹ I abbreviate ‘ $B(s)$ ’ for ‘ s is borderline’, and ‘ $D(s)$ ’ for ‘ s is definite’. Note that as I’m using the terms, ‘borderline’ and ‘definite’ as primarily predicates of sentences, or rather sentence tokens (as in ‘‘Harry is tall’ is definite’). I also allow use of these terms as sentential operators (as in ‘Definitely, Harry is tall’ or ‘It is definite that Harry is tall’). But with such uses one needs to apply similar caution as one does with propositional-attitude verbs: on the view of vagueness that is the central focus of this paper, if s_1 and s_2 are two sentences which express the same coarse-grained proposition, it is not always the case that $B(s_1)$ iff $B(s_2)$. Finally, note that some authors use the terms ‘determinately’ or ‘clearly’ (Williamson (1994)) for what I here call ‘definitely’, and some authors use the term ‘vague’ for what I here call ‘borderline’.

$B('p')$ if and only if $B('¬p')$. My focus in this paper is on a much more substantive principle which is very widely accepted²:

$$\text{Distribution: } (D('p') \wedge D('p \rightarrow q')) \rightarrow D('q')$$

Consider for example supervaluationism about vagueness. According to supervaluationism, ' p ' is definite just in case ' p ' is true on all admissible precisifications. Thus if both ' p ' and ' $p \rightarrow q$ ' are true on all precisifications, then so is ' q ', and thus ' q ' is also definite. The same is true of standard degree-theoretic views, as well as Edgington's verity theory.³

How about epistemicism? Since there is no simple analysis of the terms 'definite' and 'borderline' on the epistemicist view (more on this below), it is not trivial to determine whether or not epistemicist should accept Distribution. Nevertheless, the main proponents of epistemicism (and in particular Williamson) have explicitly endorsed the principle⁴, and I know of no other author who has denied it. Moreover, other things being equal, it seems that epistemicists should accept Distribution: for one thing the principle is intuitively plausible: it entails that definite consequences of definite statement, are themselves definite.⁵ For another, Distribution is required if we wish to treat the logic of definiteness as a normal modal logic (the principle is a version of the K axiom), which no doubt makes for a more elegant and simple logic of vagueness (e.g. it allows one to provide a Kripkean possible-world semantics for the logic).

Despite its initial appeal, I argue below that epistemicists (in particular those who follow broadly Williamsonian lines which connect ignorance due to vagueness with the phenomenon of semantic plasticity) should reject Distribution.

² The arrow here should be interpreted as the material conditional.

³ Let $V(s)$ indicate the degree of truth that s receives. On standard degree theoretic views $V('p \rightarrow q') = 1$ if $V('p') \leq V('q')$, and $1 - (V('p') - V('q'))$ otherwise. This entails that if $V('p') = V('p \rightarrow q') = 1$, then $V('p') = 1$. On Edgington's view (Edgington (1997)), $V('p \rightarrow q') \leq V('¬p') + V('q') = 1 - V('p') + V('q')$. So again, if $V('p \rightarrow q') = 1$, and $V('p') = 1$ then $V('q') = 1$.

⁴ See Williamson (1997): 271, Hawthorne (2006): 202.

⁵ Williamson justifies the principle by claiming that "if a (possibly empty) set of premises are all clear, then so are their logical consequences", (Williamson (1994): 271), but at least on one reading, the principle is even weaker than that - it only requires that if it is *definite* that something follows (in a material sense) from definite premises, then it is definite. (Cf. cautious definitions of knowledge closure which require not that knowledge be closed under entailment, but rather that it be closed under *known* entailment.)

A2. Against Distribution: the initial case

According to epistemicists there are true borderline sentences. However, even when a borderline sentence ‘*p*’ is true, we do not and cannot know that *p*. One central challenge for the view is to account for this ignorance in two related senses: first, why is it that we do not and cannot know that *p*? Second, what is distinctive about the kind of ignorance due to vagueness; how does it differ from other, more mundane cases of ignorance?

In his seminal book *Vagueness*, Williamson addresses this challenge by connecting the phenomenon of vagueness to one that we may call ‘semantic plasticity’. The idea is roughly as follows. Vague expressions such as ‘tall’ are semantically-plastic: there are close possible worlds in which speakers use ‘tall’ slightly differently than it is actually used, and this small shift in use makes for a small shift in the meaning and referent of the term.⁶ According to Williamson, the semantic-plasticity of ‘tall’ entails that, when Harry is a borderline case of tallness, one cannot know that Harry is tall (even if the claim is true). For suppose that an agent *X* were to form a true belief in the claim that Harry is tall. Then presumably there is a close possible world in which *X* would still form a belief of the same form (albeit with a slightly different content), but where due to the small shift in what ‘tall’ means in that world, that belief would be false. This in turn entails, according to Williamson, that *X*’s actual belief is not safe, and hence does not constitute knowledge.⁷

Note that while this explanation ties ignorance due to borderlineness with the phenomenon of semantic plasticity we have not yet articulated anything like a precise connection between a sentence’s being borderline and its being semantically plastic. In articulating such a connection, we can start by accepting the following two principles:

Borderlineness entails Plasticity: If *s* is borderline, then *s* contains some semantically-plastic expression.

⁶ I speak as if predicates refer to properties. One can translate the discussion into alternative semantic frameworks as required (see in particular f.n. 32 below for how my argument generalises).

⁷ Kearns & Magidor (2008) criticise this explanation of our ignorance, primarily on the basis that it appeals to an inadequate safety-principle, but I will leave that criticism aside for the purposes of this paper.

Knowledge entails Definiteness: If it is known that p then ' p ' is definite.

The first principle is justified by the thought that a necessary condition for a sentence to be borderline is that it contains a vague expression, and a necessary condition for a sentence's being vague is that it is semantically plastic. The second principle is justified by the following line of thought: suppose that it is known that p but ' p ' is not definite. By the factivity of knowledge, ' p ' is true, and ' $\neg p$ ' is false. Since the 'definitely'-operator is factive by definition $\neg D(\neg p)$, and hence $\neg D(p) \wedge \neg D(\neg p)$, which entails that $B(p)$. But the idea that can be borderline but known statements undermines the central idea of the epistemic theory of vagueness.⁸ It is also worth noting that the converse of both principles should be denied: 'Harry is bald or Harry is not bald' is definite but contains the semantically-plastic 'bald'. And clearly there are some precisely-stated mathematical truths which are definite but unknown.

With this background in place, I would like to propose a case which constitutes (I will go on to argue) a counterexample to Distribution. Let ' n ' be the canonical decimal representation of the (real) number for which the ascription "tall" refers to the property of being over n meters in height' is true (according to epistemicists, there exists such a number).⁹ I want to argue for the following three claims:

(1) *Definiteness of Semantic Cut-offs (DSC):*

D ("tall" refers to the property of being over n meters in height').

(2) *Definiteness of Disquotatation (DD):*

D ("tall" refers to being over n meters in height \rightarrow Someone is tall just in case they are more than n meters in height').

⁸ Many would take this combination to be impossible on *any* acceptable theory of vagueness, but see Dorr (2003).

⁹ As usual in these discussions, I'm idealising away from the context sensitivity of 'tall', as well as the fact that whether someone is tall might be determined by other factors than their height. I am also assuming for the purpose of the discussion that the statement 'X is more than n meters in height' is precise (if you think it is not, substitute some other precise claim that correctly specifies the truth-conditions of 'tall').

(3) *Borderlineness of Non-Semantic Cut-offs (BNC)*:

B('Someone is tall just in case they are more than n meters in height').

Clearly if we accept these three claims we have a counterexample to Distribution: if Distribution holds, then (DSC) and (DD) entail that D('Someone is tall just in case they are more than n meters in height'), which contradicts (BNC).

(DD) can be easily defended. I take it that we know – indeed know *a priori* - that if 'tall' refers to the property of being more than n meters in height, then someone is tall just in case they are more than n meters in height. So given *Knowledge entails Definiteness*, the conditional is definitely true, and we should accept Disquotational Definiteness.

(BNC) is also not particularly hard to defend. True claims of the form 'Someone is tall just in case they are more than n meters in height' seem on the face of it to be uncontroversial cases of borderline statements. One might, I suppose, object that given that I am already in the game of denying that 'tall' refers to over n meters in height' is borderline, it is no longer clear that 'Someone is tall just in case they are over n meters in height' is borderline. We need not quarrel here: let Harry be someone whose height is exactly n meters. We can now replace the three claims above with the following: D('Harry's height is n meters and 'tall' refers to being over n meters in height'); D('Harry's height is n meters and 'tall' refers to being over n meters in height \rightarrow Harry is tall')¹⁰; B('Harry is tall'). The justification of the first two claims here will carry over from my arguments in favour of the original claims ((DSC) and (DD))¹¹, and in this case, the third claim (namely, B('Harry is tall')) will be entirely uncontroversial. For ease of discussion, I will stick to the original three claims.

No doubt, though, (DSC) is the most controversial of the three claims. In a brief reply to commentators, Williamson explicitly denies it, maintaining that claims of the form 'tall' refers to the

¹⁰ I am assuming that the property of tallness is the property of having a height that is greater or equal to n meters. If it turns out that the property of 'tallness' is represented by an interval that is open in its lower end, then replace the example with one where Harry's height is n^* , where n^* is very slightly greater than n .

¹¹ This is not entirely trivial regarding DSC, but the reader can verify that this is indeed the case given the account I propose in §A.3.

property of being over n meters in height' are borderline.¹² But given *Borderlineness entails Plasticity*, this requires at least one expression in the embedded semantic ascription to be semantically-plastic. As Williamson is thinking of things, the quote-name 'tall' is not semantically-plastic (at least not in any relevant sense), and given our idealisation, 'the property of being over n meters in height' should also be treated as a precise, non-plastic term. What Williamson proposes, though, is that the word 'refers' is semantically plastic: in close possible worlds, we use 'refers' slightly differently, and this entails that a belief we would phrase in those worlds by saying "'tall' refers to the property of being over n meters in height' is false, entailing that our actual belief does not constitute knowledge.

In his 'Epistemicism and Semantic Plasticity', John Hawthorne argues against Williamson that this suggestion is untenable.¹³ The first important observation Hawthorne makes is that we ought to distinguish between semantic-plasticity and extensional-plasticity. Say that a predicate 'F' is extensionally plastic just in case in a close possible world it has a different extension than its actual extension. Extensional-plasticity does not in itself have anything to do with vagueness. Suppose I have just bet all my money in a roulette game, and (by amazing luck) won. I am in fact in the extension of the predicate 'rich'. But there is a close possible world w in which I lose all my money in the roulette game and thus do not fall under the extension of the predicate 'rich' in w .¹⁴ This extensional plasticity of 'rich' has nothing to do with vagueness. Similarly, in a close world w in which the general use facts differ slightly, 'refers' might well have a different extension than it actually has, even if 'refers' is not semantically plastic. Suppose for example that in the actual world I point to dog X and say 'Let this dog be called 'Fido'', whereas in a close world w , I utter the same stipulation while pointing to dog Y. The extension of the reference relation will vary between the two worlds (in the actual world it will contain the pair <'Fido', X>, while in w it will contain the pair <'Fido', Y>). Moreover, (even) if 'reference' picks out exactly the same relation in w as in the actual world, the extension of 'reference' will correspondingly vary between the two worlds. For exactly the

¹² Williamson (1997): 948.

¹³ Hawthorne (2006), ch. 9.

¹⁴ Of course, since 'rich' is also vague it might express a slightly different property in w , but that is also a property I clearly do not satisfy in w .

same reason, if the word ‘tall’ is used differently in a close world w , the ascription ‘‘tall’ refers to the property of being over n meters in height’ can have a different truth-value in w and in the actual world, even if no word in the ascription (and in particular not the word ‘refers’) is semantically plastic.

The upshot then is that in order to determine whether ‘refers’ is semantically plastic, it does not suffice to determine whether ‘refers’ has a different *extension* in close worlds. Rather, we need to ask whether there is a close world w in which ‘refers’ picks out a different *intension* (a function from possible worlds to extensions) than it actually does.

Hawthorne goes on to argue that the intension of ‘refers’ cannot vary in close worlds, and thus that ‘refers’ is not semantically plastic. I think Hawthorne’s argument contains a crucial flaw. In the remainder of this section I will present Hawthorne’s argument and explain why I think the argument fails. In the next section, however, I will show how one can nevertheless salvage something of Hawthorne’s insight to show that Williamson is wrong to reject *DSC*.

Here is Hawthorne’s argument for why ‘refers’ cannot be semantically plastic: suppose there is a close possible world w in which ‘refers’ expresses a different relation (i.e. a relation with a different intension) than it actually does. Call the relation ‘refers’ expresses in w ‘reference*’. Now assume that the extension of reference and reference* differ with respect to w , i.e. that there is some term ‘ n ’ and some object x , such that $\langle n, x \rangle$ is in the extension of reference but not in the extension of reference* in w . Now consider an utterance the sentence ‘‘ n ’ refers to n ’ which occurs in w . Which proposition does this utterance express? By our stipulations, in w ‘‘ n ’ refers to ‘ n ’ (the quote-name’s meaning is, we assume, stable across both worlds); ‘refers’ refers to reference*; and ‘ n ’ refers to x . Thus in w , the utterance expresses the proposition that ‘ n ’ refers* to x . But again by our stipulation, this proposition is false in w (because we are assuming that $\langle n, x \rangle$ is not in the extension of reference* in w , so the proposition that ‘ n ’ refers* to x is false in w). But the conclusion that in close worlds utterances of trivial homophonic specifications such as ‘‘ n ’ refers to n ’ are false is certainly one that Williamson

would want to reject (by his lights, it undermines the claim that in the actual world we know such specifications). Hawthorne thus concludes that ‘refers’ simply cannot be semantically plastic.

This argument contains, I think, a crucial flaw. The problem lies with Hawthorne’s assumption that if ‘refers’ has a different intension in w and in the actual world (i.e. reference and reference* have a different intension), that intension has to disagree with respect to the relevant world w . Nothing warrants this assumption:¹⁵ all that we can infer from this difference in intension is that there is *some* world v (not necessarily w itself), relative to which reference and reference* differ in extension. But this weaker assumption is insufficient for the purposes of Hawthorne’s argument. For suppose that in each close world w , ‘refers’ expresses a relation reference_w , which has the same extension as reference relative w (though possibly varies from reference relative to some other world v). Now consider an arbitrary term ‘ n ’. Suppose that in w ‘ n ’ refers to x . Since reference and reference_w have the same extension in w , then in w , ‘ n ’ also refers_w to x . Just as above, when speakers of w utter the sentence ‘‘ n ’ refers to n ’ they express the proposition that ‘ n ’ refers_w to x . But this time the proposition is *true* in w , just as we would expect.¹⁶

Where does this leave us? Recall that what we are trying to show is that the semantic cut-off claim ‘‘tall’ refers to the property of being over n meters in height’ is definite. If we restrict ourselves to the two very minimal principles governing the notion of borderlineness provided above, we have only two options: first, we can try to appeal to *knowledge entails definiteness*, arguing the ascription in question is known. But clearly, the semantic cut-off claim is not known and thus this route is hopeless. The second is to rely on (the contrapositive of) *borderlineness entails plasticity*, and argue

¹⁵ Indeed, Hawthorne’s discussion of ‘home-extension’ in §18 of Hawthorne (2006), ch. 9 suggests that he implicitly recognises this flaw.

¹⁶ Michael Caie raised the following worry: wouldn’t it be somewhat ad hoc to assume that in any close world w , ‘refers’ as used in that world happens to have exactly the same extension as the relation of reference? Response: first, I don’t think it would be ad hoc. After all, we can precisely use Hawthorne’s argument to motivate this constraint. That is, if it’s a constraint on how we use ‘reference’ that ‘‘ n ’ refers to n ’ is stably true across close worlds then we should require this congruence in extension (cf. Hawthorne’s discussion of ‘home extension’ in Hawthorne (2006), ch. 9). Second, even if we try to beef Hawthorne’s argument with the claim that the most plausible way of defending this constraint is to assume that ‘refers’ is not plastic, then his argument is at best an abductive argument, not a proof of inconsistency (as Hawthorne presents it). Finally, note of course that if we accept that ‘refers’ isn’t semantically plastic, then we can easily infer DSC (by relying on *borderlineness entails plasticity*), and thus get a failure of distribution. (For more on this point see also the discussion at the end of section A.3 below.)

that the semantic cut-off claim is definite because it is not plastic. This latter route was the one Hawthorne attempted to take, but as I argued above this attempt also fails.

It is important to remember, though, that the above principles are fairly minimal (they provide two fairly weak necessary conditions for borderliness). An alternative route is thus to beef up our theory of what it takes for a sentence to be borderline beyond these two minimal principles. In the next section, I propose a more comprehensive theory of this sort, and proceed to argue that given this theory, the semantic cut-off claim is indeed definite.

A.3 Borderliness and semantic plasticity: further connections

We have seen two minimal principles governing the notions of borderliness and definiteness within the Williamsonian framework. But can we expand these principles to a set of more informative necessary or sufficient conditions for a sentence's being borderline? As Michael Caie and others have argued, this task turns out to be surprisingly difficult.¹⁷

To simplify our task, let us restrict ourselves to the following question: what does it take for an (actually) true sentence to be (actually) borderline? The restriction to true sentences is completely benign: since a sentence is borderline if and only if its negation is borderline, we can augment our condition by adding a clause to the effect that a false sentence is borderline just in case its (true) negation is borderline. The restriction to actual-borderliness is more substantial (it is a separate question how to account for judgements such as 'If Harry were 20cm shorter, he would have been borderline tall' – it's possible that my account below can be extended to cover such sentences, but this issue will not concern me in this paper). Even with these simplifications in place, our task is far from trivial.

One natural proposal goes along the following lines¹⁸:

¹⁷ See Caie (2011), Sennet (2012), and Yli-Vakkuri (*forthcoming*).

¹⁸ Some tweaks to this characterisation are probably needed. For example, one might require that semantic plasticity is restricted to certain kinds of shift in use or that it requires some sort of indiscriminability on behalf of speakers, and the notion of a 'corresponding belief' in clause (iii) no doubt needs some fleshing out, probably

(Take-1): A true sentence s is (actually) borderline if and only if there is a close possible world w such that:

- (i) In w , due to small shifts in use, s has a slightly different meaning and expresses a slightly different proposition than it actually does (s is semantically plastic).
- (ii) The proposition expressed by s in w is such that it is false relative to w .
- (iii) If an agent in the actual world forms a (true) belief they would express using s , in w they have a corresponding (false) belief they would express using s .¹⁹

Take-1 goes a long way to accommodate the Williamsonian line of thought: clause (i) ensures that *borderlineness entails plasticity* holds while clauses (ii) and (iii) together ensure that even if an agent has a true belief in the actual world, there is a close world in which a corresponding belief is false, rendering their actual belief unsafe.

Take-1, however, is inadequate. To see why, consider the following two related problems:

Problem 1 for take-1: Suppose that in the actual world the cut-off for ‘tall’ is 1.78m and Harry’s height 1.79m. Suppose that this height is enough to ensure that Harry is not only tall, but *definitely* tall. Now, let w be a close world where the cut-off for ‘tall’ is 1.785 but Harry’s height is 1.78m. Finally, suppose that both in the actual world and in w agents have a belief that they express by saying ‘Harry is tall’ (this is easy to set up: in w Harry’s height is only just under the threshold for ‘tall’, so speakers can easily make this mistake).

by some appeal to the method of belief formation. I take it, though, that the issue of how to precisely formulate these conditions is orthogonal to the main points I make here.

¹⁹ The third clause to the effect that in w speakers have corresponding beliefs to those that they have in the actual world can alternatively be incorporated into what is meant by ‘close world’ or into what is required for semantic plasticity (clause (i)). But I prefer to write it out as a separate clause because this highlights the fact that this really is an epistemic rather than a purely semantic account of vagueness: whether a sentence is borderline depends on the doxastic attitudes of speakers, not just on the semantic and meta-semantic properties of s .

Note also that clause (iii) is phrased as a conditional in order to allow for sentences that are not actually believed to count as borderline. I am assuming here that this conditional is read as indicative, and interpreted in a way which allows it to be true even when its antecedent is false. One can alternatively restrict my account to actually true sentence that are actually believed (perhaps with some extended notion applying to sentences not actually believed), but again, I take it that these tweaks are not crucial to my argument.

It is easy to see that Take-1 predicts that ‘Harry is tall’ is actually borderline but this prediction is wrong, as Harry is definitely tall.

Problem 2 for take-1: Suppose the following claim is true in the actual world: ‘O weighs exactly 3kg and anyone who is over 2m in height is tall’. (I am assuming that the first conjunct here is perfectly precise, while the second one is vague but definite.) Let w be a close world in which O weighs 3.01kg and ‘tall’ means something slightly different than it actually does. Finally, assume that speakers in both the actual world and w have a belief that they would express by uttering the conjunction.

Again, it is easy to see that according to Take-1 the conjunction would come out as borderline (clause (i) which requires plasticity would be satisfied because the second conjunct is plastic, while clause (ii) requiring falsity in w would be satisfied because the first (precise) conjunct would be false in w). But this prediction is clearly wrong.

The problem with Take-1 is roughly this: when we are assessing whether a statement is borderline, we want to check whether we could have easily had a false belief in a case where the same underlying conditions hold – and all that has changed is the meaning of the relevant vague terms. Thus, for example, if we are assessing whether an utterance of ‘Harry is tall’ is borderline, we want to consider close worlds where Harry’s height is just as it actually is.

One proposal that has been suggested in order to address this problem is to replace clause (ii) with the following (call this move, ‘Take-2’):²⁰

(ii*) The proposition expressed by s in w is such that it is false relative **to the actual world**.

This amendment has the merit of assessing the truth of the proposition expressed in w relative to *precisely* the same underlying conditions (what better way of ensuring the conditions are the same as in the actual world than simply assessing the truth of the proposition relative to the actual world!). And indeed, this proposal delivers the right predictions in the two problem-cases described above: in the first case, Harry’s height in the actual world (1.79m) is above the cut-off for ‘tall’ in w (1.785),

²⁰ See e.g. Caie’s proposal ‘B*’, in Caie (2011).

and hence the proposition expressed by ‘Harry is tall’ in w would be true in the actual world. In the second case, the proposition expressed by the conjunction in w consists of one conjunct which is contingent but true in the actual world and another which is necessarily true – so again, the proposition would be true relative to the actual world.

But the merits of Take-2 end there. Most importantly, the amended proposal completely divorces the notion of borderliness from Williamson’s explanation of ignorance that motivated his account in the first place: after all, according to Take-2, having a true belief in a borderline statement does *not* entail that we could have easily had a false belief (the fact that we could have easily had a belief that is *actually* false is neither here nor there as far as epistemic evaluation in w goes – for this all that matters is whether one has a belief that is false *in* w). Second, and relatedly, the amended proposal simply makes the wrong predictions. Consider the following two problems.

Problem 1 for Take-2²¹: Consider the sentence ‘‘tall’ refers to property of being tall’. Since we know this claim, it must be definite. But it is far from clear that the amended proposal will deliver this prediction. Suppose that in a nearby world w , ‘tall’ refers to a slightly different property tall*. For simplicity, assume that ‘refers’ in w picks out the same relation as in the actual world.²² In w , the sentence thus expresses the proposition that ‘tall’ refers to the property of being tall*. But this proposition is false relative to the actual world (in the actual world ‘tall’ refers to the property of

²¹ See Caie (2011), §5.

²² Can this problem be avoided by supposing that ‘refers’ picks out a different relation in w ? It does not seem that it can. For suppose that in w ‘refers’ picks out a different relation – reference*. Then in w , ‘‘tall’ refers to the property of being tall’ expresses the proposition that ‘tall’ refers* to the property of being tall*. But note, first, that what Hawthorne’s argument discussed above *does* successfully establish is that reference and reference* must have the same extension relative to w (this corresponds to what Hawthorne (2006), p. 207 calls ‘domestic stability’). Thus since <‘tall’, tall*> is in the extension of reference in w , it must also be in the extension of reference* in w . But in order for the proposition expressed in w to be true relative to the actual world, <‘tall’, tall*> must also be in the extension of reference* relative to the actual world. But it would be an odd coincidence if the extension of reference* is the same in w and in the actual world (after all, the extension of *reference* isn’t stable in this way!). Moreover, as Caie points out, if such stability is true for reference*, utterances in w of counterfactuals such as ‘We could have easily referred to something different by ‘tall’’ would come out as false in w , and hence unknown in the actual world (Caie (2012), p. 375). Finally, it is worth noting that Hawthorne’s two domestic stability constraints (i.e domestic stability as it is uttered in the actual world, and as uttered in w), are not sufficient to ensure that reference* is stable in this way (To see this, consider a model where reference and reference* have a different intension but agree in extension in all nearby worlds. This model respects domestic stability, but ensures that ‘tall’ refers* to different properties in the actual world and in w because ‘tall’ refers to different properties in the two worlds.)

being tall, not to the property of being tall*). Thus assuming speakers have the relevant beliefs, the account wrongly predicts that this sentence is borderline.

Problem 2 for Take-2²³: Consider the statement ‘Harry is tall if and only if Harry is actually tall’. Since we know this statement, it has to be definite. But Take-2 predicts otherwise. For consider a close possible world w , in which ‘tall’ picks out the property of being tall*. In w , the sentence expresses the proposition that Harry is tall* if and only if Harry is tall* in w (remember that ‘actually’ is an indexical which when uttered in w picks out w). But now suppose that Harry is tall* in w but not in the actual world (recall that nothing in Take-2 requires us to consider only worlds in which Harry’s height is the same as it actually is!). This proposition would thus be false relative to the actual world (in the actual world the RHS of the bi-conditional is true but the LHS is false). Fixing the relevant beliefs of agents, Take-2 thus falsely predicts that the bi-conditional is borderline. Take-2 is not the right fix for Take-1.

Let us take stock. Take-1 was, I think, on the right track. The problem with the proposal was that we wanted to look at close worlds in which we hold fixed the relevant underlying conditions (e.g. Harry’s height in the case of ‘Harry is tall’). Take-2 attempted to fix this problem but did so in an entirely wrong way: we guaranteed that the relevant underlying facts were held fixed by assessing the truth of the proposition expressed in w relative to the actual world (instead of relative to w). But this was the wrong way to go, because it divorced the world of utterance and the world of evaluation in inappropriate ways. (Thus for example, the robustness of claims such as ‘‘tall’ refers to the property of being tall’ relies precisely on the fact that ‘tall’ is used in the same way in the world of utterance and in the world of evaluation.)

Instead, I propose that we fix the problem with Take-1 in a much more direct way, by explicitly building in the extra requirement that in the world in question the relevant underlying facts are held fixed. This, of course, leaves the tricky question of what we mean by ‘the relevant underlying facts’. In many cases, we have an intuitive grip: with respect to the claim ‘Harry is tall’, we want to keep

²³ This problem is due to Juhani Yli-Vakkuri.

fixed Harry's height; With respect to the claim 'Harry is bald', we want to keep fixed the number (and possibly arrangement) of hairs on his head. More generally, I would like to propose that given a true sentence ' p ', we need to keep fixed *the ultimate grounds* for the fact that p .

The notion of metaphysical ground has been discussed extensively in recent philosophical literature.²⁴ There are, I think, legitimate concerns about how much control or clarity we have about the concept, but hopefully, we have enough of a grip for it to be of some use here. Among those that avail themselves of grounding claims, there are a wide range of views on the grammatical structure of such claims. In particular, different authors vary on whether they treat 'grounds' as a sentential operator or as a relational predicate, and among those who take the latter view, there is disagreement about whether the arguments to which the grounding relation applies to are (possibly pluralities of) propositions, facts, or some other kinds of objects. For the purposes of the discussion below, I will assume that 'ground' is a relational predicate between facts (taking pluralities of facts in its LHS and a single fact in its RHS). However, I don't think anything rides on this assumption: the discussion can be translated into the other frameworks.²⁵

With this clarification in place, I propose, that the Williamsonian framework should adopt the following condition for borderlineness (though see §A.4 for a potential qualification)²⁶:

Condition for Borderlineness (CFB): An (actually) true sentence ' p ' is (actually) borderline if and only if there is a close possible-world w such that:

- (i) If Q are ultimate grounds for the fact that p in the actual world, then Q obtain in w .
- (ii) In w , due to small shifts in use, ' p ' has a slightly different meaning and expresses a slightly different proposition than it actually does (i.e. ' p ' is semantically plastic).

²⁴ See Correia and Schnieder (2012).

²⁵ As Fine (2012), §1.4 notes, on elaborating the view that 'ground' is a relation between facts, one would need to construe facts as sufficiently fine-grained. But at least for the purposes of my discussion below, facts only need to be as fine-grained as, roughly, Russellian propositions are (for example, for my purposes, it is not a problem if it turns out that the fact that Harry is tall is identical to the fact that Harry is over n meters in height).

²⁶ To anticipate: in §A.4, I discuss an important objection against *CFB* being a *necessary* condition for borderlineness (though one that does not threaten its status as a sufficient condition). In order to keep the discussion readable, I will continue to treat it as a necessary and sufficient condition for now, but see the discussion there for how the condition can be weakened in a way that avoids the objection but still allows my argument to go through.

- (iii) The proposition expressed by ‘*p*’ in *w* is false relative to *w*.
- (iv) If an agent in the actual world forms a (true) belief they would express using ‘*p*’, in *w* they would have a corresponding (false) belief they would expressed using ‘*p*’.

Let us see how my proposed criterion fares with respect to a range of examples:

Example 1 (‘Harry is tall’ – borderline case): Suppose that tallness is the property of being over 1.75m meters in height, and suppose that Harry’s height is 1.751m (assume that this is sufficient to ensure that ‘Harry is tall’ is borderline).

What grounds the fact that Harry is tall? Arguably, one thing that grounds the fact that Harry is tall is the fact Harry is over 1.75 meters in height.²⁷ But it does not seem that these are ultimate grounds: after all, the fact that Harry is over 1.75m in height is itself grounded in the fact that Harry is 1.751m in height. Most likely, there are in turn further things that ground this precise height fact - e.g. some micro-physical properties of Harry. But whatever exactly these ultimate grounds are, it is safe to assume that close worlds in which these ultimate grounds hold are also ones where Harry’s height is just as it actually is (namely 1.751m).²⁸

But now consider a close world *w* where the ultimate grounds for the (actual world) fact that Harry is tall are held fixed (and thus where Harry’s height is still 1.751m), but where, due to small shifts in use, ‘tall’ refers to (say) the property of being over 1.76m. In *w*, the proposition expressed by ‘Harry is tall’ would be false relative to *w*, and assuming that agents in *w* have the relevant belief, the account will correctly predict that ‘Harry is tall’ is borderline.

Example 2 (‘Harry is tall’ – definite case): A more interesting case, one which Take-1 above had trouble with, is the case of ‘Harry is tall’ when this claim is definitely true. Suppose again that tallness

²⁷ Given the irreflexivity of grounding, we should reject this claim if the fact that Harry is tall just *is* the fact that Harry is over 1.75m in height. That would not matter for my purposes, though: all I need is the claim that the fact that Harry is tall is grounded by the fact that Harry is 1.751m in height.

²⁸ There is one prominent way of thinking of ground which will easily validate this assumption, namely the view that if the fact that *q* grounds the fact that *p*, then necessarily, if *q* then *p* (see Correia (2005) and Rosen (2010)). But even if we do not accept this principle in full generality, I don’t see any reason to doubt it for this instance.

is the property of being over 1.75m, and that Harry is in fact 1.8m tall (assume this is sufficient to make ‘Harry is tall’ definite).

Now, as noted above, any world w in which the actual ultimate grounds for Harry’s being tall are held fixed is one where Harry’s height is 1.8m. But then we can assume that for close worlds w , even if ‘tall’ means something slightly different than it actually does (e.g. being over 1.76m in height), Harry would still fall above the cut-off for ‘tall’ in w and thus the proposition expressed by ‘Harry is tall’ in w would be true relative to w , ensuring that clause (iii) fails and the statement is correctly predicted to be definite.

Example 3: ‘Tall’ refers to the property of being tall’

We can now move to consider a case that caused problems for Take-2, namely the sentence ‘‘tall’ refers to the property of being tall’. As it turns out, it makes no difference what ultimate grounds here are. It is a central component of the Williamsonian apparatus that in all close worlds w , ‘‘tall’ refers to property of tall’ expresses a *true* proposition in w (this is crucial to ensuring that speakers do not violate Williamson’s safety condition in the actual world, and thus get to know that ‘tall’ refers to the property of being tall).²⁹ Thus in whatever way clause (i) restricts the relevant set of close worlds, for any such close world, clause (iii) of the condition fails and the sentence is correctly predicted to be definite.

Example 4: ‘Harry is tall if and only if Harry is actually tall’

Exactly the same applies in the case ‘Harry is tall if and only if Harry is actually tall’: since this sentence expresses a true proposition in all close worlds, my account correctly predicts that it is definite.

Example 5: ‘tallness is the property of being over n meters in height’:

²⁹ See e.g. Williamson (1997b)

Suppose that tallness is the property of being over n meters in height. What grounds this identity fact? On some ways of thinking, nothing grounds identity fact, and thus the ultimate grounds for this identity fact trivially hold in all close worlds. Others might argue that identity facts do have grounds – the only proposal I know of, maintains that the fact that $A=B$ is grounded by the fact that A exists.³⁰ But since in this case the terms that flank the identity statement are properties (and these exist necessarily), the ultimate grounds for the identity fact trivially hold in all close worlds. Either way, then, clause (i) in my account does no work in restricting the worlds under considerations.

But now we can apply clauses (ii)-(iv) as standard, and reach the conclusion that ‘tallness is the property of being over n meters in height’ is borderline (presumably, there is a nearby world where ‘tall’ is used slightly differently, and hence expresses the property of being more than n^* meters in height, and so forth...).

Example 6: ‘tall’ refers to the property of being over n meters in height’

We finally turn to the case that is the centre of my argument against Distribution: I will show that given my account of borderliness, this semantic ascription claim comes out as definite rather than borderline.

What ultimately grounds the fact that ‘tall’ refers to the property of being over n meters in height? It actually does not matter *what* exactly grounds this fact, as long as it has *some* ultimate grounds Q , and that any close worlds where these ultimate grounds hold are ones where ‘tall’ refers, just as it actually does, to the property of being over n meters in height. Surely, anyone who accepts (as Williamson does) that meaning facts are determined by use facts (broadly construed) should accept this assumption.³¹

Suppose by contradiction that the semantic cut-off sentence is borderline. Then there is a close world w satisfying clauses (i)-(iv) of *CFB*. Which proposition is expressed in w by the ascription ‘tall’

³⁰ See Burgess (2012), who argues against this proposal.

³¹ See Kearns & Magidor (2012) for a rare view that rejects this claim, but as I am assuming the Williamsonian framework here such unorthodox views are irrelevant for current purposes.

refers to the property of being over n meters in height'? The quote-name "tall" and the phrase 'the property of being over n meters in height' are by assumption precise, and thus mean in w just what they mean in the actual world. Let us suppose that 'refers' in w picks out the relation of reference* (my argument does not presuppose that 'refers' is not semantically plastic). The proposition expressed in w by the sentence is thus that 'tall' refers* to the property of being over n meters in height. By clause (iii), this proposition is false in w .

But now consider an utterance in w of "'tall' refers to the property of being tall'. Which proposition does this express in w ? The quote-name, again, does not change meaning, and 'refers', we are assuming, picks out reference*. What about 'tall' as used in w ? By clause (i), the ultimate actual grounds Q for the fact that 'tall' refers to the property of being over n meters in height hold in w , and thus in w , 'tall' must also refer (even if not refer*) to the property of being over n meters in height. Thus the proposition expressed in w by this disquotational utterance is also the proposition that tall refers* to the property of being over n meters in height. But this is precisely the proposition which we have shown to be false in w !

In summary, from the assumption that "'tall' refers to being over n meters' is borderline, we have deduced that in the relevant close world w , "'tall' refers to being tall' is false.³² But this is untenable for two reasons. First, it is extremely plausible that any ultimate grounds Q for the fact that 'tall' refers to being tall are also ultimate grounds for the fact that 'tall' refers to being over n meters in height.³³ But if this is so, then since the ultimate grounds for the latter facts hold in w , Q hold in w , and thus by our account "'tall' refers to being tall' would be borderline, which is clearly untenable.

³² It is perhaps worth highlighting that the argument does not depend on the view that predicates refer to properties, or indeed that they refer at all. In the most general terms, my argument can be put thus: let 'M' be some vague term. Let R be whatever you think the principle semantic relation is (meaning/expressing/etc. – allow it to take whatever order of term in its RHS). All I assume is that R is disquotational (i.e. "'M' R s M' is trivially true). Let 'S' be some precise term such that "'M' R s S' is a true semantic cut-off claim. If this semantic cut-off claim were borderline, there would have been a relevant close world w where this claim expresses a false proposition, i.e. where the proposition that 'M' R s S is false in w . But by clause (i), in w the actual ultimate grounds for the semantic ascription facts hold, and thus in w it is also true that 'M' R s S. So the proposition expressed in w by "'M' R s M' would be the proposition that 'M' R s S, but this proposition we have argued is false, so the disquotational claim would come out as false in close worlds which is untenable for the same reasons as above.

³³ Indeed, it is plausible that these are simply the same fact, but even if they aren't, presumably both facts are grounded by the same set of use facts.

Second, and more importantly, whether or not we get the result that ‘tall’ refers to being tall’ is *borderline* it is bad enough that it is *false* in close worlds (as noted in the discussion of Example 3 above, it is crucial for Williamson that the sentence is *true* in close worlds, for otherwise his account would predict that we cannot know this claim). The upshot, then, is that our assumption that the semantic cut-off sentence is *borderline* must be rejected, and as the sentence is true, it must be definite. DSC is thus established, and this provides a counter-example to Distribution.

Let me conclude this section by making two remarks that might help make the surprising conclusion of my argument (namely, the definiteness of the semantic cut-off claim) more palatable than it might initially appear. First, it is worth recalling that the fact that the semantic cut-off claim is definite, does not entail that it is *known* (definite claims can be unknown and indeed unknowable). Relatedly, the failure of Distribution I have argued for here need not make for a failure of knowledge-closure: the disquotational conditional is known, and has an unknown *borderline* consequent. But assuming its antecedent is a definite but unknown claim, knowledge-closure is not at risk. The claim is rather that our ignorance of the antecedent is not the distinctive kind of ignorance associated with *borderlineness*. Finally, while I was careful not to *assume* anywhere in my argument that ‘refers’ is not semantically plastic, it is worth noting that my conclusions are certainly *compatible* with the view that ‘refers’ is not plastic. Furthermore, I do think the picture on which the semantic cut-off claims are definite goes rather naturally with the view on which ‘refers’ is not plastic. On such a view, there is a simple explanation for why the semantic cut-off claim is definite: it simply contains no semantically plastic terms. Nevertheless, it is important that my argument does not rely on the assumption that ‘refers’ is not plastic.

§A.4 A potential weakening of CFB

In this final section, I discuss what I take to be the most serious objection to CFB and explain how my argument can be amended if one is swayed by this objection.³⁴

³⁴ Thanks to Cian Dorr and David Manley for raising this worry.

Here is (one version of) the objection: consider the sentence ‘The word ‘frequently’ is used frequently’. It seems that utterances of this sentence can be borderline. But at least arguably, CFB does not deliver this verdict. The problem is this: suppose that it is true that ‘frequently’ is used frequently. According to CFB, in order to assess whether this claim is borderline we ought to look only at worlds in which we hold fixed the ultimate grounds for this fact. But on the face of it, holding these grounds fixed entails holding fixed all the use facts for ‘frequently’, and this in turn would ensure that each world where the grounds are fixed would be one where ‘frequently’ has the same meaning as it actually has. But then we will not consider any worlds in which the sentence means something slightly different, and CFB would predict that the sentence is *definite* rather than borderline.³⁵

One response to this objection is to maintain that the ultimate grounds for the fact that ‘frequently’ is used frequently’ are narrower than what is needed to fix the meaning of ‘frequently’ (perhaps, for example, to fix the meaning of the word we need to consider use facts for other related words). A second response is to simply bite the bullet and claim the sentence is (surprisingly) definite.

But I would like to focus here on a third response – one which concedes that the objection undermines CFB, but offers a weakening of CFB which suffices for my purposes. For a start, note that the objection only threatens the status of CFB as a necessary condition for borderlineness, not as a sufficient condition. According to the first part of this response then, CFB should only be take to be a sufficient condition for borderlineness. Or more accurately, CFB should be replaced with the claim **CFB-S** – which replaces the ‘if and only if’ on the top of the condition with a mere ‘if’ claim. The thought is that the condition is too strong to be a necessary one: in some cases, holding fixed the ultimate grounds for the purpose of some part of the claim (e.g. how ‘frequently’ is used) does not leave sufficient flexibility with respect to another part of the sentence (namely, it does not allow the unquoted use of ‘frequently’ to be plastic).

³⁵ Of course, if ‘is used’ is vague in addition to ‘frequently’ then perhaps we won’t get the verdict that the utterance is definite. But I take it the intuition is that even if ‘is used’ were replaced by a precise predicate with the same content the utterance would be still be borderline, in virtue of the plasticity of the unquoted ‘frequently’.

This weakening on its own will not do because the crucial application of CFB in my argument above, namely its use in arguing for DSC (example 6 in §A.3 above), appealed to CFB as a *necessary condition* for borderliness. What I propose, however, is that we accept a weaker necessary (but not sufficient) condition for borderliness - call it 'CFB-N'. This weaker condition replaces clause (i) in the original CFB by the following: 'the fact that p obtains in w' (and leaves the other clauses intact). CFB-N thus relies on a very minimal sense of 'holding the relevant facts fixed': we hold the fact that p fixed, but not necessarily the ultimate grounds for this fact. (And indeed it is easy to see that this condition is too weak to be sufficient: we can hold fixed the fact that Harry is tall, but if we are not required to hold fixed the ultimate grounds for this fact we can still vary Harry's height thus bringing back the problems for Take-1). Interestingly, this weakened necessary condition is entirely sufficient to run my argument in favour of DSC: the only role that clause (i) played in the argument was in ensuring that in the close world w, 'tall' refers to over n meters (just as it does in the actual world), and this is exactly what the new condition requires. I thus propose that, in so far as one is swayed by the objection, we replace CFB (a necessary and sufficient condition) with two weaker principles CFB-S and CFB-N (one which is sufficient but not necessary and one which is necessary but not sufficient).

A final note: given that my argument for DSC only requires CFB-N one might wonder whether CFB-S is doing any interesting explanatory work. The answer is positive: first, my objections to Take-1 and Take-2 were raised primarily against the *sufficiency* of these conditions, so providing an alternative sufficient condition that does not face these problems is highly desirable. Second, and more importantly, the counter-examples to Distribution I discuss in part B (see esp. §B.4) crucially relies on CFB as *sufficient* condition.³⁶ While not providing a necessary and sufficient condition, together CFB-S and CFB-N provide an interesting and informative theory of borderliness which makes for various kinds of failures of Distribution.

³⁶ The counterexamples there also require some necessary condition, but only an extremely weak one as all the definiteness claims defended involve either entirely non-plastic terms, or at sentences where one does not have any false beliefs in *any* close worlds. Certainly, CFB-N more than suffices for those purposes.

Part B: Epistemicism and the argument from vagueness

§B.1 The argument from vagueness

Universalism about composition states that every class of objects has a fusion. Arguably the most influential argument for universalism is the ‘argument from vagueness’, first suggested by David Lewis and then substantially developed by Ted Sider.³⁷ The argument has been widely discussed and criticised.³⁸ I am sympathetic to some of these criticisms, but I will not discuss them here. Rather, my aim in this second part of the paper is to address the following question: assuming the argument is otherwise successful, do epistemicist have any particular reason to reject it?³⁹

Sider’s argument relies on three premises which he argues for in turn. Here they are verbatim:

P1: If not every class has a fusion, then there must be a pair of cases connected by a continuous series of cases such that in one, composition occurs, but in the other, composition does not occur.

P2: In no continuous series is there a sharp-cut off in whether composition occurs.

³⁷ See Lewis (1986): 212-213, Sider (1997): 214-22, and Sider (2001): 120-132.

³⁸ See Korman (2010) for an overview of this debate.

³⁹ The question of whether epistemicists have any particular reason to reject the argument from vagueness is only discussed briefly (and rather unsatisfactorily) in the current literature. Hudson (2000) assumes the answer is obviously positive, but his only comment on why this is so is that “the epistemic theory of vagueness... would threaten to undermine the defences of Lewis and Sider by blocking the inference from a vague restriction on composition to the repugnant conclusion that existence is possibly indeterminate” (*ibid.* 549). I take it that Hudson is trying to resist P3, though he does not elaborate on how this line would proceed.

Sider (2001): 130-2 discusses the question briefly. He claims that while it might be thought the epistemicists can reject P2, they should not ultimately reject this premise (his remarks are in line with my comments in §B.2). He does not, however, mention any reason for the epistemicist to resist P3.

López De Sa (2006): 403 briefly argues that epistemicist cannot resist the argument by denying that ‘exists’ (understood unrestrictedly) is precise. I do not contest this claim.

Korman (2010), n. 16 makes the following remark: “One who opts for an epistemicist strategy [for rejecting the argument] will need to find a way to resist A4 [‘there cannot be cut-offs with respect to composition’], but will also presumably reject A3 [‘there cannot be borderline cases of composition’] given the usual epistemicist gloss on ‘borderline’ and ‘determinate’”. I am not sure neither why, according to Korman, A3 (which corresponds to P3 above) would be denied given an epistemicist understanding of ‘borderline’, nor why he says that even if the epistemicist resists A3, they must also resist A4. At any rate, his brief and tentative suggestion for how the epistemicist might resist P3 is by allowing a sentence to be borderline even if no expression in it is vague. The strategy I offer in §B.3 does not rely on accepting this highly problematic commitment.

Finally, Steen (2014), assumes that the epistemicist would obviously reject P2 (but offers an alternative argument to Sider’s for why the epistemicist should accept unrestricted composition).

P3: In any case of composition, either composition definitely occurs or composition definitely does not occur.

Taken together the three premises entail that the antecedent of P1 ought to be rejected: that is, every class of objects has a fusion.

One might be tempted to think that the epistemicist can easily respond to the argument by rejecting P2. In §B.2 below, I explain why this quick response is misguided: epistemicists do not have any particular reason to reject P2. In §B.3 and §B.4, however, I argue that the rejection of Distribution argued for in Part A does provide epistemicists with grounds for resisting P3, or more accurately for rejecting Sider's argument in support of this premise. The upshot is that epistemicists do have a reason to reject the argument from vagueness. The discussion also provides a further potential counterexample to Distribution (and interestingly, one that – unlike the counterexample presented in Part A - does not involve any semantic terms).

§B.2 Why epistemicists have no particular reason to reject P2

It might be tempting to assume that epistemicist can easily reject P2: after all, don't epistemicists accept that even obviously vague predicates such as 'bald' or 'tall' have sharp cut-offs? The problem with this response is that it is based on an equivocation of the term 'sharp cut-off'. No doubt there is a common use of the term in the vagueness literature according to which a predicate has a 'sharp cut-off' when, roughly, its uses are subject to the laws of excluded middle and of bivalence. And of course on this common reading epistemicists are indeed committed to the claim that even 'tall' or 'bald' have sharp cut-offs.⁴⁰

But this is not the way Sider uses the term. As he clarifies: "By a 'sharp cut-off' in a continuous series I mean a pair of adjacent cases in a continuous series such that in one, composition definitely occurs, but in the other, composition definitely does not occur".⁴¹ Crucially, this interpretation can remain

⁴⁰ See e.g. Williamson (1994): 13 for this use.

⁴¹ Sider (2001): 123. It is somewhat distracting that Sider later shifts between the two uses. Thus for example he says: "Since epistemicists are already accustomed to accepting sharp cut-offs for 'heap' and 'bald', one might

neutral on the interpretation of the ‘definitely’ operator, and in particular allow for an epistemicist-friendly understanding of the term.

Not only does this second reading conform to Sider’s explicit elaboration of P2, but I think it is by far the more charitable reading of his argument. First, Sider himself insists that epistemicists should not deny P2, so surely we should prefer a reading of the premise which is consistent with epistemicism. Second, this is the reading that is required to make Sider’s argument *valid*, as P3 is phrased in terms of definite composition rather than in terms of ‘sharp cut-offs’. Third, Sider’s defence of P2 (both in relation to epistemicism and to his preferred semantic views) does not rely in any way on the issue of rejections of bivalence or LEM (as the other reading of ‘sharp cut-offs’ would predict). His claim is rather that both semantic and epistemicist views of vagueness rely on there being a range of equally eligible candidate meanings for a term (thus, for example, there is nothing in our metaphysics to privilege one potential precisification of ‘tall’ over another). He maintains that what distinguishes epistemicism is the commitment to the claim that our use facts succeed in picking out a single candidate meaning, *even though* that meaning is not metaphysically privileged.⁴² Finally he maintains that no two adjacent cases in a sorites series for ‘compose’ will be such that one exemplifies a more metaphysically privileged relation. His argument (particularly the claim that no candidate ‘composition’ relation is more eligible) is certainly not uncontroversial, but reasons to reject it have nothing in particular to do with epistemicism.⁴³ There is nothing to bar epistemicists from agreeing that, just as with ‘tall’, there are no two adjacent cases in the relevant sorites series such that in one way have definite composition and in the other definite lack of.

think they would also be happy with a sharp cut-off in a continuous series of cases of composition, thus rejecting P2...I will argue that even an epistemicist should accept P2” (*ibid.* 131). I take it that what Sider should have said is that in the sense of ‘sharp cut-off’ relevant to P2, epistemicist would not accept such cut-offs for ‘heap’ or bald’ any more than they would for ‘compose’.

⁴² This interpretation of epistemicism seems well in line with Williamson’s remarks on the connection between lack of metaphysical privilege and semantic plasticity (see Williamson (1994): 231).

⁴³ Hawthorne, for example, objects to this line on the grounds that even if it seems that there is no metaphysically substantive differences when looking at the series from a micro-physical perspective, there might be ones from a “higher-level” emergent perspective (See Hawthorne (2006):104-109)) but this point has nothing in particular to do with epistemicism.

Even putting aside Sider's own intentions, the above remark suffice to show that there is interesting and apparently compelling version of his argument on which both P2 and P3 involve a neutral notion of 'definitely'.⁴⁴ On that version of the argument, epistemicists need not reject P2.

§B.3 Why epistemicists do have a reason to resist P3

No doubt the most innovative part of Sider's argument consists of his argument in favour of P3. The argument proceeds roughly as follows: in so far as there are violations of P3, there are violations in settings involving only finitely many objects. But in such settings, borderliness about whether composition occurs makes for borderliness about how many objects there are. The problem is that in finite settings, questions about how many objects there are can be phrased in purely logical terms. And since no logical term is vague, claims about how many objects there are not vague, and hence cannot be borderline.

Now before I discuss the argument in more detail, it is important to note that there is no *prima facie* reason why this argument would not be acceptable to epistemicists: the crux of the argument relies on a tension between the claim that logical terminology is not vague, and the claim that (assuming a rejection of universalism) 'compose' is vague. But both of these are claims that, at least on the face of it, epistemicists have as much reason as anyone to accept.⁴⁵

Sider's own discussion of the argument for P3 focuses almost entirely on defending the claim that (in finite settings) questions about how many objects there are cannot be vague. I will not contest this claim (and indeed, I do not think epistemicists have any special reason to reject it). Rather, I would like to focus on a conditional which Sider assumes with very little discussion: (in finite settings), if it is borderline whether a certain class of objects compose, then there is some corresponding claim about how many objects there are which is also borderline.

⁴⁴ Of course, I do go on to argue that on this reading epistemicists do have the resources to reject P3, but only for very delicate reason involving failures of Distribution. As I explain below, there is certainly no *obvious* reason for epistemicists to reject P3 or Sider's argument in its favour.

⁴⁵ See López De Sa (2006) for an explanation of why Sider's argument that the relevant logical terminology is precise should be equally acceptable to epistemicists.

How might this claim be supported? A starting thought is that claims about whether a certain class of objects compose entail claims about how many objects there are. Suppose for example, that we start out with two atoms. One might be tempted to accept the following two conditionals:

Composition Conditional 1 (CC1): If the two atoms compose, then there are at least three objects.

Composition Conditional 2 (CC2): If the two atoms do not compose, then there are fewer than three objects.

One problem is that (CC2) is far too quick: suppose for example that the domain contains exactly three objects: the two atoms and a third object which is either a simple co-located with the two atoms or else an object composed of the two atoms. Whether or not the two atoms compose, there are three objects in the domain.⁴⁶ But let us leave these reservations aside: if nothing else, they again have nothing in particular to do with epistemicism (in so far as the argument can be resisted in this way, this solution is open to any theorist of vagueness). Assume, then, that we are in a setting where both (CC1) and (CC2) hold, and where the following also holds:

(B) ‘The two atoms compose’ is borderline.

Now consider the proposed connection between the borderlineness of the composition and the borderlineness of the count-claim:

Conditional Borderlineness (CB): If ‘the two atoms compose’ is borderline then ‘There are at least three objects’ is borderline.

Are the assumptions (CC1) and (CC2) sufficient to establish (CB)? It is far from trivial to show that they are. For one thing, on no plausible theory of vagueness can we infer from ‘*p* is borderline’ and ‘If *p* then *q*’ that ‘*q* is borderline. For example, we certainly shouldn’t accept that the claims that ‘Harry is bald’ is borderline and that if Harry is bald then Harry has fewer than one million hairs,

⁴⁶ See Korman (2010), §4, Hawley (2004): 389, and Magidor (forthcoming) §3.2.2 for related objections to this part of Sider’s argument.

together entail that ‘Harry has fewer than one million hairs’ is borderline. The situation here is, of course, somewhat different because (CC1) and (CC2) together entail a bi-conditional, and it might be tempting to think that if one side of a true bi-conditional is borderline then so is the other. But this principle, too, would certainly be rejected by epistemicists (even putting aside my rejection of Distribution). For example, epistemicists accept that there are true bi-conditionals of the form ‘Harry is tall if and only if Harry’s height is greater than n meters’, where the LHS is borderline, and the RHS is definite.

Let me then propose on Sider’s behalf, a more carefully stated argument.⁴⁷ Suppose that in the relevant settings where (CC1) and (CC2) hold we also *know* that they hold and thus know (using some basic logic) that if there are at least three objects the atoms compose, and if there aren’t at least three objects the atoms do not compose. Since (given the epistemic theory) we have *knowledge entails definiteness*, we can conclude the following two conditionals:

(DCC1): ‘If there aren’t at least three objects, the two atoms do not compose’ is definite.

(DCC2): ‘If there are at least three objects, then the two atoms compose’ is definite.

Now, assume that we also accept Distribution, then from DCC1 we can get that *if $\neg D(\neg$ the atoms compose) then $\neg D(\neg$ there are at least three objects)* and from DCC2 that *if $\neg D$ (the atoms compose) then $\neg D$ (there are at least three objects)*. But together, these claims entail that *if B (the atoms compose) then B (there are at least three objects)*, namely (CB). Together with (B) this would entail that B (there are at least three objects), a conclusion which – Sider maintains – we cannot accept.

To summarise: while stating Sider’s argument properly requires some care, if we accept Distribution (as well as some background assumptions about the settings), an enhanced version of Sider’s argument in favour of P3 would be just as effective for epistemicists. The problem, of course, is that this enhanced argument crucially relies on Distribution. Thus given my argument in Part A, epistemicists should not accept the above argument in favour of P3.

⁴⁷ Cf. Hawthorne’s reconstruction of the argument in Hawthorne (2006): 105.

One might have the following lingering worry: even if epistemicists do not accept Distribution in general, they still accept that some particular instances of the principle are correct. Perhaps then the particular instances appealed to in the above argument are unproblematic?

Once the general principle of Distribution is rejected, the burden of proof for showing that the instances used in the argument are correct lies with Sider. Nevertheless, to strengthen my argument, in the next section I present two models on which the relevant instances of Distribution appealed to in Sider's argument fail. These will serve both to show that the instances of Distribution appealed to in the argument are far from innocuous, as well as to provide some further potential counterexamples to Distribution. Moreover, as opposed to the counterexample presented in Part A, these potential examples do not involve any semantic claims. The latter point is important: if there was any suspicion that the original counterexample had only arisen because funny things go on when one applies a theory of vagueness to semantic claims – these kinds of non-semantic examples serve to show that the problem is much more widespread.

§B.4 Failures of Distribution for Sider's case

In this final section, I provide two models on which DCC1, DCC2, and B hold, while CB fails. My discussion will assume the condition for borderliness *CFB* defended in §A.3 (though as I note in the discussion, the weaker condition defended in §A.4, namely the conjunction of *CFB-S* and *CFB-N*, equally supports my argument). Naturally, the models involve quite a bit of simplification and idealisation and require various controversial assumptions to set up. It is worth clarifying that I do not take these models to provide anything like a conclusive proof that the particular instances of Distribution that Sider appeals to in fact fail. I do, however, take them to show how these instances *might well* fail. This puts further pressure on Sider's burden of proof, as well as demonstrating why failures of Distribution need not be restricted to sentences involving semantic claims.

Model 1 (CC1 and CC2 are true in all nearby worlds):

Let us start with a general underlying picture about how the meaning of a term is determined by use facts. There are a wide range of use facts for a given term. Many of these involve sentences speakers typically assent to, though of course there is a large variety in factors such as how often speakers assent to these, how confident they are in their assent, as well as how expert are those assenting.⁴⁸ As a rough picture we can assume that these use facts place constraints that *other things being equal* the meaning of the term should respect. Of course, often there is simply no candidate meaning that respects *all the constraints*, and in that case the meaning will be determined according to some ‘best fit’ with use, where ‘best fit’ can take into account some sort of ranked list of how central different constraints are to our use practice.

Let us assume the actual world ($w_{@}$) contains exactly two objects: two atoms, A and B, located exactly 1.05mm apart.⁴⁹ All close worlds also contain A and B, and do not contain any objects which are located elsewhere. Let us suppose that the use facts for the word ‘compose’, both in the actual world and in close worlds, place the following constraint that the meaning of ‘compose’ should (other things being equal) respect, listed here in descending order from most to least important:

- (1) A and B compose if and only if there are at least three objects.

Comment: this corresponds to accepting Sider’s two conditionals CC1 and CC2.

- (2) A and B compose if and only if they are sufficiently close to each other.

Comment: I am assuming that there is some vagueness about what we mean by ‘sufficiently close’, but let us assume the only two legitimate precisifications are either ‘closer than 1mm’ or ‘closer than 1.1mm’.

- (3) *Necessarily*, A and B compose if and only if they are sufficiently close to each other.

Comment: I am assuming that speakers endorse not only claim (2) but also its necessitation.

However, they might be somewhat less confident about the modal claim, and hence it ranks below claim (2).

⁴⁸ Of course, other kinds of use facts involve general features of the world (for this broad interpretation of the term ‘use facts’ see Williamson (1994) §7.5).

⁴⁹ This is, of course, a simplification as presumably we want other objects in the domain such as speakers. But it is clear how the simplified picture can be amended here (e.g. let the world contain an additional number of n objects, and then shift all count claims accordingly).

(4) *Necessarily*, if A and B compose, then there are at least 3 objects.

(5) *Necessarily*, if A and B do not compose, then there are fewer than 3 objects.

Comment: Claims (4) and (5) are the necessitation of CC1 and CC2 respectively. I am assuming here that speakers not only endorse CC1 and CC2, but also their necessitation. However, they might be somewhat less confident in these modal claims, hence these rank below claim (1). I am also assuming that the necessitation of CC1 ranks a bit higher than the necessitation of CC2 (after all, it seems particularly compelling that presence of composition is a *count increasing* matter).

Now consider a close world w_1 , which contains exactly three objects: A and B which are located exactly as they are in the actual world (and in particular are 1.05mm apart), but also contain an additional object C, which is co-located where A and B are (jointly) located.

What is the intension of ‘compose’ in $w_@$? I propose that given the above ranked list of constraints, a very good candidate for the intension of ‘compose’ in $w_@$ is one where, necessarily, A and B compose if and only if they are less than 1mm apart. This would validate not only the most important constraint (1), but will in fact allow us to validate each of constraints (1)-(4). The only constraint that would be violated here is constraint (5) (as on this interpretation in w_1 the two atoms do not compose but there are more than 3 objects).

What is the intension of ‘compose’ in w_1 ? I propose, that given the above ranked list of constraints, a good fit is some intension on which, ‘A and B compose’ is true *relative to* w_1 if and only if A and B are less than 1.1mm apart (and thus, as it turns out, ‘A and B compose’ is true relative to w_1). This will also us at the very least to maintain the two most important constraints (1) and (2).⁵⁰

⁵⁰ Will the intension satisfy the other constraints? It is not trivial to ensure that it satisfies constraint (3). For let *compose** be the relation that ‘compose’ expresses in w_1 . Satisfying (3) would require that A and B *compose** in $w_@$, even though it seems as if there is no object which they *compose**. This can perhaps be accommodated by various non-standard metaphysical views. For example, one might insist that A and B *compose** A (for a defence of possibilities of this sort see Kearns (2011) and Bacon and Cotnoir (2012)). Or we might relax the connection between composition* and location, and insist that A and B *compose** some third object located elsewhere (adjusting the count claims accordingly). Or one might adopt a more restricted version of the deflationary view described below, so that A and B can *compose** even though there is no (single) object they *compose**.

With this set-up in place, we can now see why this model would provide a relevant counter-example to Distribution (i.e. one where DC1, DC2, B hold, but CB fails):

Why do (DCC1) and (DCC2) hold on this model? Because constraint (1) is the most high-ranking constraint both in the actual world and in close worlds. This ensures that in each close world w , whatever ‘compose’ means in w , the propositions expressed in w by CC1 and CC2 will be true, thus ensuring that clause (iii) of *CFB* fails.⁵¹

Why does (B) hold on this model? I will show this by arguing that in $w_{@}$ ‘A and B do not compose’ is borderline true. Recall that in $w_{@}$ the intension of ‘compose’ is such that ‘Necessarily, A and B compose iff they are less than 1mm apart’ is true. For a start, this ensures that in $w_{@}$, ‘A and B do not compose’ is *true* (because A and B are 1.05mm apart). Next, we need to determine what *grounds* the fact that the two atoms do not (actually) compose. Given that the intension of ‘compose’ is such that whether the two atoms compose is determined entirely by the distance between them, then very plausibly, what grounds the fact that the two atoms do not compose is the fact that they are 1.05mm apart. (Just as what grounds the fact that Harry is tall is his precise height.).

Now consider w_1 . Since in w_1 the atoms are still exactly 1.05mm apart, the ultimate (actual) grounds for the (actual) fact that the two atoms do not compose holds in w_1 , as required by clause (i) of *CFB*. Next, in w_1 ‘compose’ has a different intension than it has in $w_{@}$: we know the intensions must be different because they differ in their extensions relative to w_1 : the actual intension makes ‘A and B do not compose’ true relative w_1 , whereas the w_1 -intention makes the sentence false relative to w_1 . This ensures that clause (ii) of *CFB* holds. Next, as we have just noted, the proposition expressed in w_1 by ‘A and B do not compose’ is false in w_1 , ensuring clause (iii) of the condition holds. And finally, we can simply stipulate that agents’ doxastic attitudes are such as to make clause (iv) hold. This concludes the proof that (B) holds in this model.

But more plausibly the ‘best fit’ for ‘compose’ in w_1 simply does not succeed in satisfying constraint (3). (In that case, (4) and (5) can still be satisfied). At any rate, note that my argument does not depend in any way on how compose* behaves anywhere other than w_1 so we do not need to settle this issue.

⁵¹ And obviously, it would fail even if we adopt the weaker necessary condition CFB-N.

Finally, why does (CB) fail on this model? Simply because the claim ‘There are at fewer than 3 objects’ is (I’m assuming, following Sider) perfectly precise, so by clause (ii) in the condition it cannot be borderline. This ensures that while ‘A and B do not compose’ is borderline, the corresponding count claim is *not borderline* and thus we get the relevant failure of Distribution.

Some might feel uneasiness about the fact that $w_{@}$ and w_1 contain a different number of objects. If one follows the argument, one should realise that this difference does not undermine the force of the example. Nevertheless, the next model is one on which both of the worlds in question are ones that contain exactly the same number of atoms (arranged in precisely the same way).

Model 2: (CC1 is false in a nearby world, but not believed)

To set up this model, we will start by considering two very different conceptions of composition. The first I call ‘SUBSTANTIVE’. According to SUBSTANTIVE, necessarily, atoms compose just in case they are sufficiently close to each other (for simplicity, assume that in the case of two atoms, they are required to be closer than 1mm to each other). The second conception we’ll call ‘DEFLATIONARY’. According to DEFLATIONARY, at a *semantic* level, composition is a many-to-many relation (i.e. a relation between pluralities), even if syntactically ‘compose’ behaves as if it were a many-to-one relation.⁵² Moreover, on this conception necessarily, a plurality of objects compose another plurality just in case they are *the very same plurality*: the *xxs* compose the *yys* if and only if the *xxs* are the *yys*.⁵³ According to DEFLATIONARY, any plurality of objects trivially composes (as any plurality of objects are identical to themselves). Note, however, that numerical statements still count only individual objects rather than pluralities. According to DEFLATIONARY, then, two atoms A and B, necessarily compose (because necessarily, they compose themselves), even though the fact that they compose does not entail that there are any more than two objects.

⁵² See Cotnoir (2013) for a defence of the claim that the syntactic and semantic features of a relational predicate can diverge in this way.

⁵³ One can thus think of DEFLATIONARY as an unorthodox version of ‘composition as identity’ (unorthodox because most proponents of the view take composition, as well as identity, to be many-to-one in these cases).

Now suppose that both in the actual world and in close worlds use facts for ‘compose’ are such that the only two potential meanings for the term are DEFLATIONARY and SUBSTANTIVE. How can it be that two such radically different conceptions are both candidate meanings for the term? Suppose that speakers feel some pull towards each of the two conceptions. In some moods, they feel that composition must be substantive and assent to sentences such as ‘Objects compose just when they are sufficiently close to each other’. In other moods, they feel that composition must be deflationary and utter sentences like ‘There is no more to saying that the atoms compose than saying that there are some atoms’. Thus the use facts pull in both direction, and which of the two meanings ‘compose’ gets assigned depends on the fine-grained details of which of the two conceptions gets slightly favoured in the use facts. Let us further assume the following:

- (1) All close worlds, including $w_{@}$, contain exactly two objects: atoms A and B which are one mile apart.⁵⁴
- (2) In $w_{@}$, the use facts slightly favour SUBSTANTIVE, entailing that this is what ‘compose’ means in $w_{@}$.
- (3) In some close worlds, the use facts slightly favour DEFLATIONARY, entailing that this is what ‘compose’ means in those worlds. As noted above, the deflationary conception is not friendly to CC1. Let us assume, then, that part of the way that DEFLATIONARY is slightly favoured in the relevant worlds is that speakers in those worlds do not assent to CC1 (though they might be hesitant to outright reject it). Let w_1 be one such world which favour DEFLATIONARY.
- (4) Suppose that both in $w_{@}$ and in w_1 speakers assent to ‘A and B do not compose’ (in both worlds this pattern is part of their pull towards the SUBSTANTIVE, though only in $w_{@}$ does this pull ultimately win).

Why do (DCC1) and (DCC2) hold on this model? This is because even though utterances of CC1 are false in some close worlds (e.g. w_1) the pattern of the use facts ensures that in all close worlds in which this sentence is false, speakers do not have beliefs that they would express using (CC1),

⁵⁴ Similar qualification as in n. 49 apply here.

so clause (iv) for the *CFB* (and *CFB-N*) fails. (*DCC2* holds trivially, because the consequent of *CC2* is true in all nearby worlds).

Why does (B) hold on this model? I will show this by arguing that in $w_{@}$ ‘A and B do not compose’ is borderline. Recall that the intension of ‘A and B do not compose’ in $w_{@}$ is true relative to a world w just in case A and B are closer than 1mm in w . For a start, this ensures that ‘A and B do not compose’ is true in w_1 (after all, they are one mile apart). Next, we need to determine what *grounds* the fact that the two atoms do not (actually) compose. Since ‘compose’ in $w_{@}$ means ‘less than 1mm apart’ what grounds the fact that the two atoms do not compose is that they are one mile apart. Now consider w_1 . Since in w_1 the atoms are also exactly 1 mile apart, the ultimate (actual) grounds for the (actual) fact that the two atoms do not compose hold in w_1 , as required by clause (i) of *CFB*. Next, in w_1 ‘compose’ has a different intension than it has in $w_{@}$: we know the intensions must be different, because they differ in their extensions relative to w_1 : the actual intension makes ‘A and B do not compose’ true relative w_1 , whereas the w_1 -intension makes it false relative to w_1 . This ensures that clause (ii) of the condition holds. Next, as we have just noted, the proposition expressed in w_1 by ‘A and B do not compose’ is false in w_1 , ensuring clause (iii) of the condition holds. And finally, given that we have stipulated that both in $w_{@}$ and w_1 speakers have beliefs that they would express by saying ‘The two atoms do not compose’, clause (iv) holds. This concludes the proof that (B) holds in this model.

Finally, why does (CB) fail on this model? Simply because the claim ‘There are fewer than 3 objects’ is (I’m assuming, following Sider) perfectly precise, so by clause (ii) in the condition it cannot be borderline. This ensures that while ‘A and B do not compose’ is borderline, the corresponding count claim is not borderline, and thus we get the relevant failure of Distribution.

Model 2 thus gives a further example of how the relevant failures of distribution might occur.⁵⁵

⁵⁵ For helpful comments and discussion, thanks to Michael Caie, Cian Dorr, Jeremy Goodman, John Hawthorne, David Manley, Timothy Williamson, Juhani Yli-Vakkuri, audiences in Birmingham, Norwich, Oslo, Princeton, and St. Andrews, as well as the anonymous reviewers.

References

- Bacon, A. and Cotnoir, A. (2012): 'Non-wellfounded mereology', *Review of Symbolic Logic* 5: 187-204.
- Burgess, A. (2012): 'A puzzle about identity', *Thought* 1: 90-99.
- Caie, M. (2012): 'Vagueness and semantic indiscernibility', *Philosophical Studies* 106: 365-377.
- Correia, F. (2005): *Existential Dependence and Cognate Notions*, Munich: Philosophia.
- Correia, F., and Schnieder, B., eds. (2012): *Metaphysical Grounding: Understanding the Structure of Reality*, Cambridge: Cambridge University Press.
- Cotnoir, A. (2013) 'Composition as general identity', *Oxford Studies in Metaphysics* 8: 295-322.
- Dorr, C. (2003): 'Vagueness without ignorance', *Philosophical Perspective* 17: 83-113.
- Edgington, D. (1997): 'Vagueness by degrees', in R. Keefe and P. Smith (eds.), *Vagueness: A reader*, MIT Press: 294-316.
- Fine, K. (2012): 'A guide to ground', in Correia & Schnieder (2012): 37-80.
- Hawley, K. (2004): 'Borderline simple or extremely simple', *The Monist* 87: 385-404.
- Hawthorne, J. (2006): *Metaphysical Essays*, Oxford: Oxford University Press.
- Hudson, H. (2000): 'Universalism, four dimensionalism, and vagueness', *Philosophy and Phenomenological Research* 60: 547-60.
- Kearns, S. (2011): 'Can a thing be a part of itself?', *American Philosophical Quarterly* 48: 87-96.
- Kearns, S. & Magidor, O. (2008): 'Epistemicism about vagueness and meta-linguistic safety', *Philosophical Perspectives* 22: 277-304.
- Kearns, S. & Magidor, O. (2012): 'Semantic Sovereignty', *Philosophy and Phenomenological Research*: 322-350.
- Korman, D. (2010): 'The argument from vagueness', *Philosophy Compass* 5: 891-901.

- Lewis, D. (1986): *On the Plurality of Worlds*, Oxford: Basil Blackwell.
- López De Sa, D., (2006): ‘Is ‘everything’ precise?’, *Dialectica* 60: 397-409.
- Magidor, O. (*forthcoming*): ‘Endurantism vs. perdurantism?: a debate reconsidered’, *forthcoming* in *Nous*.
- Rosen, G. (2010): ‘Metaphysical dependence: grounding and reduction’, in B. Hale and A. Hoffmann, *Modality: Metaphysics, Logic, and Epistemology*, Oxford: Oxford University Press: 109-136.
- Steen, I. (2014): ‘Almost Ontology: Why epistemicism cannot help us avoid unrestricted composition or diachronic vagueness’, *Pacific Philosophical Quarterly* 95: 130-139.
- Williamson, T. (1994): *Vagueness*, London: Routledge.
- Williamson (1997): ‘Replies to commentators’, *Philosophy and Phenomenological Research* 57: 945-53.
- Williamson, T. (1997b): ‘Reply to commentators: [Horgan, Gomez-Torrente, Tye]’, *Philosophical Issues* 8: 255–256.
- Yli-Vakkuri, J. (*forthcoming*) ‘Epistemicism and modality’, *forthcoming* in *Canadian Journal of Philosophy*.