Scope Confusions and Unsatisfiable Disjuncts:
Two Problems for Supervaluationism

DELIA GRAFF FARA
Princeton University

Abstract

Here we elaborate two problems for supervaluationist accounts of vagueness. (I) The best (canonical-)supervaluationist explanation of our inclination to accept sorites premises attributes to us a tendency to confuse the scopes of a *Truth* operator with the existential quantifier. This explanation is shown to be incorrect as well as incomplete. (II) A well-known complaint against supervaluation semantics is that it allows for a disjunction to be true even though none of its disjuncts is in fact true. Here we develop a new, related complaint: supervaluation semantics allows for a disjunction to be true even though none of its disjuncts *could* be true.

I. Introduction

Supervaluationism as a theory of vagueness has its advantages and its disadvantages. Most of the advantages of supervaluationism are ones that favor it over those of its rivals that also reject bivalence. Since I believe in bivalence, those advantages do not hold much sway over me. I think that the disadvantages of supervaluationism far outweigh its advantages. But I offer no cost–benefit analysis here. Rather, I want to provide some detailed discussion of a couple of the disadvantages.

The view discussed here is canonical supervaluationism, which I'll take to be the view presented by Kit Fine in his 1975 article and defended by Rosanna Keefe in her more recent book (Keefe 2000).[1] On this view, a claim is *supervaluationally true* (or just *true*) when it is true-according-to-classical-semantics on all of the different admissible ways of 'precisifying' the vague expressions in its language.[2] The same goes for falsity: a claim is *false* (*simpliciter*) when it is false on all of the different admissible ways of collectively precisifying the vague expressions in its language. Since some claims involving vague expressions are true on some admissible precisifications but false on others, there are claims containing vague expressions that are deemed by the supervaluationist to be neither true nor false. This is what it is to reject bivalence.

A predicate is vague on the canonical view only if its extension gap is non-empty, only, that is, if there are objects of which it is neither true nor false. To precisify a vague predicate is to assign all of the objects in its extension gap to one or the other of its extension (the things of which it's true) or anti-extension (the things of which it's false). What makes some precisifications of the vague expressions jointly admissible?[3] Those precisifications must not alter the truth

[1] But see also Hans Kamp (1975) and Dominic Hyde (1997) for very different applications of the supervaluationist techniques to some problems of vagueness.

[2] By a *claim*, I mean an utterance with the following properties: it's made in a particular language, it's made in a particular context, and it says something (it 'expresses a proposition'). When I say that a claim *contains a certain expression*, I mean that the sentence of which that claim is an utterance contains that expression.

[3] The precisifications must be done collectively, since some admissible precisification of 'small', for example, will render a number of precisifications of 'tiny' inadmissible: everything tiny is small, and there are some small things that aren't tiny; in semantic talk, everything in the extension of 'tiny' on a precisification must be in the extension of 'small' on that precisification.

value of any claim that already has a truth value. Some of these claims will relate different vague predicates to one another (‘no-one huge is tiny’); others will relate vague predicates to associated relational expressions (‘everyone larger than someone huge is also huge); still others will relate vague predicates to their clear cases (‘anyone under one meter tall is tiny’).[4] These claims are said to express ‘penumbral connections’—ones that hold even when the quantifiers in them range over borderline cases of the vague predicates in them: over the things in the predicate’s ‘penumbral region’.[5] It would be mistaken to describe these constraints by saying that the supervaluationist deems some precisifications admissible whenever they’re compatible with our current usage of vague terms, since on her view a lack of extension gap is not compatible with vagueness. (This turns out to be the underlying source of the focus of §3, below)

After a brief inventory in §2 of four advantages and disadvantages of canonical supervaluationism, we will focus on two of the disadvantages: (§3) canonical supervaluationism allows not merely for true disjunctions with no true disjuncts, but also for true disjunctions with no *satisfiable* disjuncts; (§4) supervaluationists have yet to provide us with any convincing answer to the question of how we could ever find the false premise of a sorites argument to be so appealing.

II. Benefits and Costs

The main advantages of supervaluational treatments of vagueness are these:

[4]Which cases are clear cases may vary with context. If ‘tiny’ means ‘tiny for an adult’, then one meter is a clear case; though not if it means ‘tiny for a two-year old’.

[5]The terminology is adapted from Bertrand Russell (1923).

a. Their account—in terms of truth-value gaps—of what it is to be a borderline case of a vague predicate is satisfying to most philosophers not wedded to bivalence (which unfortunately may well be most philosophers).

b. They preserve the truth of truisms expressing penumbral connections, such as ‘anyone shorter than a short person is short’ or ‘everything tiny is small’, whereas rival theories that also ditch bivalence do not.

c. They account for the fallacy involved in soritical reasoning: in deeming *false* the inductive premise of a sorites argument (*e.g.*, ‘any man just one nanometer taller than a short man is himself also short’), they declare the argument unsound.

d. They preserve classical logic, for the most part anyway.[6]

Not surprisingly, these advantages come with their disadvantages. (No theory is perfect.) The chief ones of these are:

a. a commitment, *despite* the rejection of bivalence, to vague predicates’ having sharp boundaries—in the sense of there being a least tall height, one nanometer below which renders a person not tall,[7] a reddest red, a thinnest fat person, etc;

[6]See Timothy Williamson (1994) and Delia Fara (2004, §2) for arguments that the preservation of classical logic is incomplete in important ways.

[7]This isn’t exactly accurate; the commitment is rather to there being either a shortest tall height *or* a tallest non-tall height, and likewise for the other cases mentioned. When variation occurs along a continuous scale, a region can be bounded without itself containing that boundary.

b. a rejection of standard forms of reasoning, such as contraposition of arguments and *reductio ad absurdum*; and

c. a failure to provide a satisfactory explanation of why we're mistakenly tempted by sorites reasoning.

d. an assignment of extremely counter-intuitive truth conditions to the logical particles other than negation;

III. True Disjunctions with Unsatisfiable Disjuncts

Supervaluation yields truth conditions that are extremely counter-intuitive for the logical particles other than negation. It is often criticized, and justly so, for allowing there to be true disjunctions or true existential generalizations that have no true disjunct or instance; and correspondingly, false conjunctions or false universal generalizations that have no false conjunct or instance. These failures of the classical truth conditions are alleged to arise due to what Fine called a 'truth-value shift'. In the case of disjunctions and existentials, different disjuncts or instances can be the verifying ones on different admissible precisifications; but as long as there is at least one verifier on any given admissible precisification, the disjunction or the existential will be true *simpliciter*. The main problem with this as a justification, as I will discuss shortly, is that a sentence can be true on an admissible precisification without its even being *possible* for the sentence to be true *simpliciter*. This makes room for there being true disjunctions both of whose disjuncts could not possibly be true. That there could be a shift in truth from disjunct to disjunct when neither disjunct *could* be true is repellent.

I say: 'Someone in this room is the shortest tall person'. Supervaluationists say: 'We agree, but there needn't be a correct answer to the question *Who?*'. And I say: 'Either Juan or Carlos is the shortest tall person in this room'. Supervaluationists say: 'We agree, but there needn't be a correct answer to the question *Which one?*'.

And supervaluation allows all statements made in the following dialog to be correct:

> YOUR CHAIR: Don't be late for the 4:00 meeting.
>
> YOU: Ok, what's the latest I can be without being late?
>
> CHAIR: A short time after the hour.
>
> YOU: I figured that, but *what* time?
>
> CHAIR: I couldn't truly tell you any particular time.

This leaves me staring incredulously, as it should you. In fact, it's such a strange position to hold that one often hears this said of supervaluationists:

> They think that (i) there's a shortest tall man, even though they don't think that (ii) someone in particular is such that *he* is the shortest tall man.

But of course, supervaluationists do believe both of these things; like all of us, they regard these claims as trivially equivalent. What they do believe, which the mistaken attributer is trying to capture, is that there is a shortest tall man, even though no-one in particular is such that it's *true* that he is the shortest tall man. Supervaluationism is so discordant with the way we actually speak that there are philosophers who understand the view but who aren't yet fluent in the language we would be speaking *if* it were correct. But these philosophers *are* fluent in the language we're in fact speaking. Whatever prescriptive merits supervaluationism might have, it is not descriptively correct.

The following supervaluationist explanation of the anomaly has significant appeal.

> The reason 'It's either pink or red' can be true even when neither disjunct is true results from the vagueness of 'pink' and 'red', and in particular from the fact that the object in question is a borderline case of both predicates. We can tell by looking at the thing that it falls somewhere on the spectrum between pink and red, which explains why the disjunction is true. But it's indeterminate which of pink or red it is, since there are different ways of drawing a boundary between pink and red that count as admissible precisifications of our vague usage of these words, and the object falls on different sides of that boundary on different ones of these ways. And that is why neither disjunct is true *simpliciter*.

The problem with the explanation is that it does not sit at all well with the following fact about supervaluational semantics. It allows not only for true disjunctions where neither disjunct is true, it allows also for true disjunctions where neither disjunct could be true. I assume that the supervaluationist must say that a sentence *could not be true* when there is no supervaluational model in which it is true.

Let us use $\Diamond_s$ in the following way: $\Diamond_s \phi$ is true on a precisification just in case Φ is true in some supervaluational model.[8] $\Diamond_s$ is the supervaluational *satisfiability* operator. Then for an appropriately chosen Φ and Ψ, supervaluationists allow for the truth of the following:

$$(\Phi \vee \Psi) \wedge \neg \Diamond_s \Phi \wedge \neg \Diamond_s \Psi.$$

The disjunction of this appropriately chosen Φ and Ψ is true because at least one of them is true on every precisification, even though each disjunct is truth valueless, and even though neither disjunct *could* be true. If we are to accept the claim of the form '$\Phi \vee \Psi$, but it's indeterminate which' on the grounds that things could *go either way* depending on how you drew precise boundaries for the predicates involved, then it should be indeed that things *could* go either way. But in the case to be presented, it is not.

---

[8]The connection with genuine possibility should be apparent. One would think that possibility should amount to truth in some possible world, where on the supervaluationist's conception each possible world would correspond to some supervaluational model, namely that model that verified and falsified exactly the same sentences as it. Depending on the supervaluationist's treatment of modality, the diamond used here might not be the diamond of modal logic, since it remains open for the supervaluationist to say that for the modal operator $\Diamond$, $\Diamond\Phi$ is true if Φ is true in some possible world, but false only if Φ is false in every possible world. In that case, since our Φ (and Ψ too) will be untrue in every supervaluational model, but also neither true nor false in some supervaluational models, $\Diamond\Phi$ would be neither true nor false, while $\Diamond_s\Phi$ is false.

So what's the example? It's very simple. Suppose we have something that's a borderline case of 'pink'. Let's represent the claim that it is pink as p. This claim is indeterminate—i.e., neither true nor false according to supervaluationism. We'll use B to stand for the borderline-case operator. So Bp is true. Supervaluation verifies every instance of excluded middle, even for indeterminate propositions like p. Take the true disjunction $p \vee \neg p$; conjoin each of its disjuncts with the true claim Bp; then the result, $(Bp \wedge p) \vee (Bp \wedge \neg p)$, is true.[9] But neither of its disjuncts could be true.

To see that neither could be true, note that when Bp is true, neither p nor $\neg p$ can be true: we'll have p true on some admissible precisifications of 'pink', $\neg p$ on others. But Bp manages to be true on all of these precisifications, for its truth at a precisification depends on the values p takes over the whole space of precisifications: Bp is true on a precisification when p is true on some admissible precisifications, but false on others of them.[10] This is not something that can vary from precisification to precisification, so Bp is true *simpliciter*.

But Bp is *incompatible* with p, and also with $\neg p$. For the truth of Bp requires precisely that neither of these claims be true on every precisification, that each be true on some, false on others. So, we have the supervaluational truth of the

[9] Another way to see the commitment is to note that the conjunction of two true claims is true. So $Bp \wedge (p \vee \neg p)$ is true. By classical distribution laws, $(Bp \wedge p) \vee (Bp \wedge \neg p)$ is true.

[10] The the truth conditions associated with the B operator are structurally like those associated with contingency in modal logic. Those for the B operator involve quantification over admissible precisifications, while those for contingency quantify over possible worlds.

following, when Bp is true:

$$((Bp \wedge p) \vee (Bp \wedge \neg p)) \ \wedge \ \neg\Diamond_s(Bp \wedge p) \ \wedge \ \neg\Diamond_s(Bp \wedge \neg p).$$

We have substituted $(Bp \wedge p)$ for Φ and $(Bp \wedge \neg p)$ for Ψ in the schema given earlier.

The underlying source of the difficulty is that supervaluationists (i) think that the existence of borderline cases is not compatible with precision, understood as no gap in truth value; yet they (ii) supply a semantics on which claims of borderline status remain true upon complete precisification of the vague expressions in the language. They are constrained to supply such a semantics by their thoughts that (iii) there are claims of borderline status that are true *simpliciter* while (iv) truth *simpliciter* is to be identified with truth-on-every-precisification.[11] Putting these views together requires there to be claims of borderline status (ones of the form $B\Phi$) that are true on each precisification, even though it must be that either Φ or its negation is true on that precisification. I take (i–iv) to be at least partly constitutive of of the view I'm here criticizing; to reject any of these in light of my criticism is to concede defeat. In particular, to give up (i), the incompatibility of precision with the existence of borderline cases, is to allow for bivalence in spite of vagueness; while to refuse ii) is to reject at least one of (iii)

[11] *Cf.* Keefe (2000, 186f.). Here Keefe discusses whether supervaluationists face the problem that many of their own theoretical claims (*e.g.* 'a predicate is vague only if it has an extension gap') are not true on any precisification, much less on all of them. She argues that the best supervaluational response is to supply a semantics of the kind described as feature (ii) above. On that point, I agree with her.

and (iv). To reject (iii) is to reject vagueness. To reject (iv), meanwhile, is to reject supervaluationism at its core.

I would like to say that if a sentence *cannot* be true, then that suffices for its being impossible. The supervaluationist might object to this philosophical platitude, but if I were to help myself to it, I could put the main point of this section this way. Supervaluationists think that a single impossibility need not be false, and that the disjunction of two impossibilities can be true. This is because they think that it's impossible for something to be both borderline pink and pink.[12] Yet, for anything that is a borderline case of pink, the claim that it's both borderline pink and pink can only be indeterminate, hence not false, according to their neither-true-nor-false account of indeterminacy.

IV. Scope Confusion and the Psychological Question

Supervaluation falsifies the inductive premises of sorites arguments, but does not directly provide an answer to the question of why we're mistakenly inclined to believe them in the first place. Kamp was prompted to spurn his (1975) supervaluational theory of vagueness for just this reason. Fine, who published an independently-developed supervaluational account of vagueness that same year, did offer some explanation (1975, 286). I have argued elsewhere that that explanation was unsatisfactory (2000, 50–52). This was primarily on the grounds that it implicitly required that we tend to equate truth with non-falsity, something which

[12]This is not the case on the typical view that affirms bivalence, although Raffman's (2005) pro-bivalence view is a salient exception.

an opponent of bivalence cannot do without undermining his own theory. Keefe, one of the most prominent current supervaluationists, has offered a supervaluationist answer to this ‘psychological question,’ as I have called it (2000, 50).

The psychological question is the question why we are so inclined to believe a sentence with the form of (U) in many cases where it is in fact false. I know of only two places where a supervaluationist answer to the psychological question is offered or defended: in Fine’s (1975, 286) article and in Keefe’s (2000, 183–186) book.

(U) $\forall x(\Phi x \rightarrow \Phi x')$.

Here Φ is some sorites-susceptible predicate and x' is the successor of x on some sorites series for that predicate.

By a *supervaluationist* answer to the question, I mean one that derives from supervaluationism *per se*, not from supervaluationism as it might be supplemented with some further view better designed to answer the psychological question, such as a contextualist theory or some boundary-shifting theory more generally.[13] The main flaw in both of Fine’s and Keefe’s explanations is an implicit reliance on our tending to equate truth with non-falsity. This is at best in serious tension with the supervaluationist rejection of bivalence. In the remainder of this section I argue against the success of Keefe’s answer in particular to the psychological question.[14]

[13]Boundary-shifting theories, whether they be contextualist ((Kamp 1981), (Raffman 1994, Raffman 1996), (Soames 1999)) or invariantist (Fara 2000, Fara 2008) are not in competition with any particular account of what it is to be a borderline case of a vague predicate, where these latter accounts include epistemicism, degree theories, supervaluationism, or truth-value-gap theories more generally.

[14]I have argued against Fine’s elsewhere (Fara 2000, 50–52)

Keefe's scope-confusion explanation of our mistake

Keefe's explanation of why we believe (U)—when we do, which as she points out (183f), may not be always—'turns on the fact that supervaluationism can distinguish between the falsity of (U) and its having a false instance, and correspondingly between (E)'s being true and its having a true instance' (184). ((E) is an equivalent of the negation of (U).)

(E) $\exists x(\Phi x \wedge \neg \Phi x')$,

(TE) $\text{TRUE} : \exists x(\Phi x \wedge \neg \Phi x')$,

(ET) $\exists x\ \text{TRUE} : (\Phi x \wedge \neg \Phi x')$.

'TRUE:' here represents the operator 'it is true that'. We mistakenly confuse (ET) with (TE) (and hence with (E), an equivalent of the latter), so that our correct denial of (ET) mutates into an incorrect denial of (E).

The problem with the explanation

The first question to address is how exactly a confusion of (ET) with (E), and a concomitant confusion of the denial of (ET) with the denial of (E), explains our incorrect attitude towards an entirely different sentence, (U). It behooves us to spell out the chain of reasoning. We begin with the correct denial of (ET), which, due to scope confusion, leads to an incorrect denial of (TE). Then a chain of good logical reasoning leads us to the incorrect affirmation of (U). The connecting links are spelled out below. In the 'Reason' column, I cite the justification for the judgement made in that row on the basis of the judgement made in the immediately preceding row. In the 'Correctness' column I indicate whether the *judgement* made in that row is good, not whether the reasoning is good. Each step of reasoning is good, except for the scope confusion made at the outset.

Table 1.

	Sentence	Judgement	Reason	Correct?
(ET)	$\exists x$ TRUE : $(Fx \wedge \neg Fx')$	Deny	Good Judgement	Correct
(TE)	TRUE : $\exists x(Fx \wedge \neg Fx')$	Deny	Scope Confusion	Incorrect
(3)	TRUE : $\neg\forall x(Fx \rightarrow Fx')$	Deny	Substitution of Equivalents	Incorrect
(4)	FALSE : $\forall x(Fx \rightarrow Fx')$	Deny	FALSE : $\Phi \equiv$ TRUE : $\neg\Phi$	Incorrect
(5)	TRUE : $\forall x(Fx \rightarrow Fx')$	Affirm	????	Incorrect
(6)	$\forall x(Fx \rightarrow Fx')$	Affirm	True : $\Phi \equiv \Phi$	Incorrect

What's missing, as I've indicated, is an explanation of why we move from step (4) to step (5). Why would we go from denying a certain falsity ascription to affirming the *truth* of its embedded clause? For given the supervaluationist rejection of bivalence, there are two ways for a claim to fail to be false. It could be true or it could be indefinite.

Two further questions are pressing. First, why is our good judgement directed at (ET) rather than (TE); why don't we correctly affirm (TE) and *then* because of scope confusion incorrectly affirm (ET), leading us eventually to the correct conclusion that the inductive premise is false?[15] If anything, our attitude to (TE) should be the dominant one, since of the two sentences, only it (according to the supervaluationist) is equivalent to a simple non-metalinguistic claim.

Second, why do we tend toward confusion of the relative scopes of a quantifier and a truth-value operator only in the case of $\exists$ and 'TRUE:'? If we ignore quantifiers other than the ones typically appearing in formal first-order languages, there are four possible scope confusions to be considered. The following table represents the only plausible position for the supervaluationist to take on the chance of our making any of these scope confusions. It is telling that the only case in which we'd be at all inclined to confuse the relative scopes of a truth-value operator and a quantifier is the one where the supervaluational opinion about equivalence differs from the classical opinion.

Table 2.

Equivalence?	Chance for Mistake?	Classical?	Supervaluational?
$\exists x\ \text{TRUE} : \Phi x \stackrel{?}{\equiv} \text{TRUE} : \exists x \Phi x$	Yes	Yes	No
$\exists x\ \text{FALSE} : \Phi x \stackrel{?}{\equiv} \text{FALSE} : \exists x \Phi x$	No	No	No
$\forall x\ \text{FALSE} : \Phi x \stackrel{?}{\equiv} \text{FALSE} : \forall x \Phi x$	No	No	No
$\forall x\ \text{TRUE} : \Phi x \stackrel{?}{\equiv} \text{TRUE} : \forall x \Phi x$	No	Yes	Yes

[15]I have in mind this chain of inference: Correct affirmation of $\text{TRUE} : \exists x\ (\Phi x \wedge \neg \Phi x') \Longrightarrow$ Incorrect affirmation of $\exists x\ \text{TRUE} : (\Phi x \wedge \neg \Phi x') \Longrightarrow$ incorrect affirmation of $\exists x\ \text{TRUE} : \neg(\Phi x \rightarrow \Phi x') \Longrightarrow$ incorrect affirmation of $\exists x\ \text{FALSE} : (\Phi x \rightarrow \Phi x') \Longrightarrow$ correct affirmation of $\text{FALSE} : \forall x (\Phi x \rightarrow \Phi x')$.

The table displays, incidentally, that something like the reasoning in Table 1. must be involved in our inference to the wrong conclusion in the case of (U)—if that inference is to involve a scope confusion. There is no *other* illegitimate scope confusion of the relevant kind which we're in danger of making.

Keefe does allege that that we are in danger of confusing the relative scopes of 'TRUE:' and ∃ in at least one other case. I don't find the example convincing. She writes:

> [The confusion of (E) and (ET)] is thus like a confusion between saying that it is true that someone ought to do X and saying that it is true of someone that they ought to do X: the latter may be false while the former is true. We would run the two together if we thought the only way that 'someone ought to do X' could be true was if there was someone, y who ought to do X. But ... the former could hold because X being done is a right of z's and so it ought to be done by someone, though it is no individual's *duty* to do it (185).

In my assessment, we should use a deontic operator 'OUGHT:' to represent the conflated sentences as follows.

(OE) TRUE: OUGHT: $\exists y X y$,

(EO) $\exists y$ TRUE: OUGHT: $X y$.

There is a scope confusion alright. But it is one between OUGHT: and ∃, not TRUE: and ∃, and can be represented without involving truth at all:

(OE*) OUGHT: $\exists y X y$,

(EO*) $\exists y$ OUGHT: $X y$.

Let me summarize the points made in this section. The scope-confusion explanation of our failure to recognize (U) as false fails for the following reasons: first, the scope explanation succeeds only if we're apt to confuse non-falsity with truth, which we should not be if supervaluational semantics were correct; second, scope confusion is symmetric, but the explanation on offer requires an unexplained and unlikely asymmetry in the inferential order of our judgements; third, there is no good explanation for why we might make such a scope confusion, since we're not at all in danger of doing so in any of the relevantly similar cases (Table 2); while fourth, the only supposedly clear example of the confusion in question is in fact an example of an unrelated sort of scope confusion.

V. Conclusion

No plausible or satisfying supervaluationist answer to the psychological question has yet been offered. I probably phrased the question somewhat badly at the outset, however. But it is to the question so phrased that Keefe, and Fine before her, answered it. The problem with the phrasing is that it demands an explanation of our attitude to a false *universal generalization*, whereas what is really needed is an explanation of why we're inclined to believe each *instance* of the generalization given that we have overwhelming evidence that not all of its instances are true. Without such an explanation, we have no answer to the question why we're inclined to accept sorites reasoning when its premises do not include a universal generalization but rather, in the place of one, a series of its instances—a series of claims about adjacent pairs of minimally differing objects, *e.g.*, 'these are either both tall or both not' or 'this one is tall if that one is'. Combined with this, ideally, would be an explanation of why we're unable to locate a shift in *any* kind of status along a sorites series, not only from truth to falsity, but also, e.g., from truth to truth-valuelessness, from clear cases to borderline cases, or even from clear cases to cases about which nothing relevant could truly be said (perhaps, not even that).

A supervaluationist might remedy the deficit by supplementing her view with some complementary contextualist or boundary-shifting answer to these questions. The considerations in §3, however, suggest that we should be leery of the prospects for ultimate success. Since supervaluationists allow for true disjunctions with only unsatisfiable disjuncts, they cannot legitimately appeal to 'truth-value shift' in order to explain away the strangeness of the truth conditions they assign to disjunctions, conjunctions, etc.

References

Fara, Delia Graff (2000), 'Shifting Sands: An Interest-Relative Theory of Vagueness', *Philosophical Topics* **28**(1): 45–81. Published under the name 'Delia Graff'.

Fara, Delia Graff (2004), 'Gap Principles, Penumbral Consequence and Infinitely

Higher-Order Vagueness', in J. C. Beall, ed., *Liars and Heaps: New Essays on Paradox*, Oxford University Press, pp. 195–221. Published under the name 'Delia Graff'.

Fara, Delia Graff (2008), 'Profiling Interest Relativity', *Analysis* **68**(300): 326–335.

Fine, Kit (1975), 'Vagueness, Truth and Logic', *Synthese* **30**: 265–300.

Hyde, Dominic (1997), 'From Heaps and Gaps to Heaps of Gluts', *Mind* **106**(424): 641–660.

James, E., Slater, J. et al., eds. (1983–), *The Collected Papers of Bertrand Russell*, Allen & Unwin/Unwin Hyman, London.

Kamp, Hans (1975), 'Two Theories about Adjectives', in E. L. Keenan, ed., *Formal Semantics of Natural Language*, Cambridge University Press, Cambridge, England, pp. 123–155.

Kamp, Hans (1981), 'The Paradox of the Heap', in U. Mönnich, ed., *Aspects of Philosophical Logic*, D. Reidel, Dordrecht, pp. 225–277.

Keefe, Rosanna (2000), *Theories of Vagueness*, Cambridge University Press, Cambridge, England.

Raffman, Diana (1994), 'Vagueness without Paradox', *Philosophical Review* **103**(1): 41–74.

Raffman, Diana (1996), ‘Vagueness and Context-Relativity’, *Philosophical Studies* **81**: 175–92.

Raffman, Diana (2005), ‘Borderline Cases and Bivalence’, *Philosophical Review* **114**(1): 1–31.

Russell, Bertrand (1923), ‘Vagueness’, *Australasian Journal of Philosophy and Psychology* **1**: 84–92. Page references are to reprint in E. James, J. Slater et al., eds. (1983–).

Soames, Scott (1999), *Understanding Truth*, Oxford University Press, New York.

Williamson, Timothy (1994), *Vagueness*, Routledge, London.