Vagueness, contextualism, and ellipsis*

Paul Elbourne
University of Oxford

Abstract Stanley (2003) has argued that contextualist theories of vagueness are inconsistent with a certain fact about the interpretation of indexicals in Verb Phrase ellipsis, namely that the semantic content of an indexical in an elided verb phrase must be the same as the semantic content of the corresponding indexical in the antecedent verb phrase. In this paper, some counterexamples are adduced to undermine confidence in this generalization and hence Stanley's argument as a whole.

Keywords: vagueness, sorites, ellipsis, indexicality, contextualism

1 Introduction

Stanley (2003) put forward a criticism of contextualism about vagueness that is still influential.[1] This being the case, and in spite of the fact that Stanley's argument has already received a certain amount of critical attention (Ellis 2004, Raffman 2005, Gert 2008), this article launches a new criticism against it. Briefly, it will argue that the empirical generalization upon which Stanley relies is falsified by counterexamples.

Contextualist accounts of the sorites claim that vague terms are context-sensitive. In particular, indexicalist contextualist accounts claim that this context-sensitivity consists in the semantic content of vague terms (and thus

---

1 For example, it is the only criticism of this kind of contextualism described in detail in the survey article on the sorites by Sainsbury & Williamson (2017). The contextualist theory of vagueness was originated by Kamp (1981); other prominent contextualists include Raffman (1994), Soames (1999), Fara (2000), and Shapiro (2006).

their extension) being different on different occasions (Soames 2002: p. 445); this seems to be the kind of context-sensitivity that Stanley (2003: pp. 270–1) has in mind.[2] The idea is that when we are presented with two adjacent items in a sorites sequence, provided that they are sufficiently similar, we will unconsciously adjust the extension of the vague term so that both fall under it, even if we would not have been inclined at first to say that the more peripheral of them fell under the term if it had been presented in isolation. This is supposed to explain the naturalness of sorites arguments, despite their deleterious conclusions.[3]

Vague terms, then, on the kind of account that we are examining, are indexicals.[4] This is where Stanley comes in. He claims the following (2003: p. 271):

(1) *Invariant Interpretation of Indexicals (III)*
     Indexicals have invariant interpretations in Verb Phrase (VP) ellipsis.

What is meant can be understood from the following examples of VP ellipsis, due to Stanley:

(2) a. John likes me and Bill does too.
     b. Hannah lives here and Bill does too.
     c. John read that and Bill did too.

As Stanley points out, (2a) cannot be used to claim that Bill likes anyone other than the speaker of the first conjunct, even if the second conjunct is

---

2 Some authors make a distinction, in discussions of vagueness, between this kind of context-sensitivity and so-called non-indexical contextualism, which claims that the context-sensitivity of vague terms leaves the semantic content alone and affects their extensions by manipulating the circumstances relative to which they are to be evaluated. See Åkerman 2012: pp. 471–2 for further discussion. I will not be dealing with non-indexical contextualism in this article.

3 Different contextualist authors have had different means of bringing about this unconscious adjustment of the extension of vague terms. Kamp (1981: pp. 237–8, 243), for example, posits the psychological influence of a principle he calls Equivalence of Observationally Indistinguishable entities (EOI), which runs as follows (where P is a vague predicate): 'Suppose the objects a and b are observationally indistinguishable in the respects relevant to P; then either a and b both satisfy P or else neither of them does.'

4 It is perhaps important to clarify that all I mean by this term, and all Stanley seems to mean given his use of *that* in example (2c), is an item whose content is dependent on the context of use (and whose content can thus differ on different occasions). This is also how Kaplan (1989: p. 490) understood the term in his classic paper on the subject. I do not have in mind any narrower category like the pure indexicals (Kaplan 1989: p. 491).

uttered by someone else. (2b) has to be used to say that Bill lives in the same place that Hannah does. And (2c) has to be used to say that John and Bill read the same thing. Examples can be multiplied. A plausible generalization is that the indexicals in question have to refer to the same things (or, more generally, have the same semantic content) in the antecedent verb phrase and the elided verb phrase in each case, which is what Stanley means in his generalization.

Here is how Stanley (2003: p. 272) thinks this principle is problematic for indexicalist contextualist theories of the sorites. Imagine we have a row of n agglomerations of grains. The first member of the series is clearly a heap. Each successive member contains one grain fewer than the last until, when we reach the nth member, we are clearly not dealing with a heap. (Or so we might think at first, if that agglomeration of grains is presented in isolation.) Suppose someone points at each agglomeration in turn and says the following, where each $that_i$ is used to demonstrate the ith member of the series:

(3)   If $that_1$ is a heap, then $that_2$ is too, and if $that_2$ is, then $that_3$ is, and if $that_3$ is, then $that_4$ is,…, and if $that_{n-1}$ is, then $that_n$ is.

Now the contextualist, says Stanley, wishes to say that the reason each conditional in this utterance would be compelling in context is that, whenever we consider an adjacent pair of agglomerations of grains, we tacitly adjust the content of *heap* so that its extension includes both. But this tactic cannot work with this example, he claims, since each occurrence of *heap* after the first occurs in an elided verb phrase; by III, this means that each occurrence of this word in fact has the same content (the content provided by the first, overt, occurrence), contrary to what is required by the indexicalist contextualist account of the sorites.

Thus Stanley (2003). But I will now argue that III is false.

2   Some counterexamples to III

An example that seems to be inconsistent with III involves the indexical *current*. Very frequently, this word seems to contribute the time of utterance to the proposition expressed in a context-sensitive manner.[5] If spoken at 3pm

---

[5] I will consider other uses of the word, and whether they undermine the counterexample being presented here, in Section 3.2.

on 28 March 2019, for example, (4) means that the British government in power at 3pm on 28 March 2019 is in danger of falling:

(4)   The current British government is in danger of falling.

Against this background, and in light of III, it comes as a surprise, then, to process the following example:

(5)   Sarah (the conservative) likes the current government and...*(speaking very slowly, while a sudden coup takes place)* Mary (the revolutionary communist) does now.

Suppose that the overt occurrence of *current* is spoken at precisely 3pm on 28 March 2019 and *does* (which immediately precedes the elided verb phrase) is spoken precisely one minute later, following a very rapid and efficient communist putsch.[6] Then there is a prominent (and perhaps the only) reading of the example whereby it is used to assert that Sarah likes the government in power at 3pm on 28 March 2019 and Mary likes the government in power at 3.01pm on the same day. This is obviously inconsistent with III.

Here is a neat variant of example (5), suggested by an anonymous reviewer for *Semantics & Pragmatics*:

(6)   *(A commentator on a NASCAR race starts to mention the current leader at the very moment when another driver makes a move and overtakes him.)* Jones is the current leader and...well, wait, now Smith is!

Given these examples, it is now open to the indexicalist contextualist to claim that vague words like *heap* are also indexicals that fail to comply with III.

Examples (5) and (6) can be of assistance to indexicalist contextualists in another way too. Åkerman (2012: pp. 476–7) (crediting Elia Zardini) has drawn attention to the following variant of Stanley's version of the sorites paradox:

(7)   The first guy in the line is tall relative to the current standings of the relevant contextual factors, and if the first guy is, then the second guy is too, and if the second guy is, then the third guy is too,..., and if the $n - 1$th guy is, then the nth guy is too.

---

6 The exact times do not matter, however, and the interval could be made shorter if having such a long one seems suspicious. Of course, the example by Kaplan (1989: p. 514) on which this one is loosely based had the relevant interval be that between the appearance of the Morning Star and the appearance of the Evening Star.

The idea here is that some indexicalist contextualists might claim that the kind of indexicality associated with vagueness might not be subject to III, even though common-or-garden indexicality is. (In fact Gert (2008) has produced a proposal along these lines.) The variant in (7) is meant to show that the sorites is still a problem even if one eradicates vagueness by fixing contextual factors in a way that contextualists might be expected to approve of. By means of the qualification *relative to the current standings of the relevant contextual factors*, we are supposed to fix the interpretation of *tall* once and for all to some determinate value; the idea is then that we can still construct a sorites sequence, thus showing that the sorites paradox cannot be caused by shifting context. But it is doubtful that this example does what it is meant to. To start with, the exact workings of the relevant contextual factors are still to a large extent mysterious; and their setting and unsetting is a process that operates beneath conscious awareness. It is extremely dubious, then, to assume that we can set them manually, as it were, by means of explicitly referring to them in the manner attempted in (7). But even if we overlook this problem, we see that the use of the indexical *current* does not, contrary to intention, ensure that one and only one value for the relevant contextual factors is in play throughout the example, since, as we see in (5) and (6), this word can itself take on different values over time, even across antecedent and elided verb phrase.

3   Complications and objections

3.1   Other possible counterexamples

I am not averse, of course, to the idea that there are other possible counterexamples to III. And, indeed, a quick inspection of the literature on ellipsis will reveal some plausible candidates. But I think it is more difficult than it might seem at first sight to establish that these candidates really are counterexamples.

A convenient place to start is with the strict/sloppy ambiguity (Ross 1967). (8a) is ambiguous between (8b) and (8c):

(8)   a.   John loves his mother and Bill does too.
      b.   '…Bill loves Bill's mother.' (Sloppy)
      c.   '…Bill loves John's mother.' (Strict)

The sloppy reading might seem on the face of it to be a counterexample to III; but probably no researcher in the field would cite this as a counterexample, since a well-established and possibly correct explanation of the sloppy reading is to say that it derives from the pronoun *his* being bound (Keenan 1971). Perhaps we are dealing with two different LFs, along the following lines (where the index 1 is mapped to John):

(9) *Sloppy*
     John [λ_2 t_2 love his$_2$ mother] and Bill [λ_2 t_2 love his$_2$ mother]

(10) *Strict*
     John [λ_2 t_2 loves his$_1$ mother] and Bill [λ_2 t_2 love his$_1$ mother]

This kind of solution, of course, would claim that the antecedent and elided VPs had identical semantic content in the sloppy case, something like 'loves x's mother' (informally speaking), with the 'x' bound by a higher operator each time. Thus the sloppy case would not be a counterexample to III.

Things are different, however, when we come to examine the following examples, in which the crucial pronouns are not in the scope of any potential binders:

(11) The policeman who arrested John read him his rights, but the policeman who arrested Bill didn't. (Wescoat 1989)

(12) If John has trouble at school, I'll help him, but if Bill does, I won't. (Hardt 1999)

Example (11), for example, clearly has a reading whereby it claims that the policeman who arrested Bill did not read him, Bill, his rights; the claim is not that the policeman who arrested Bill did not read John his rights. And if I say (12), I am expressing an intention not to help Bill if he, Bill, has trouble at school; I am not expressing an intention not to help John if Bill has trouble at school. And yet the pronouns in the ellipsis sites in these examples cannot be bound by the relevant potential antecedent (*Bill*) in either case, since this word is embedded deep within islands in both cases and cannot plausibly have scope over the pronouns in question; it does not seem possible to extend the previous analysis to these cases, then.

This kind of example has consequently provoked at least one analysis that challenges III: Fiengo & May (1994) claim that examples like (11) and (12) work by employing sets of referential indices in isomorphic patterns in antecedent and ellipsis sentences. If this is on the right lines, then these

examples already constitute counterexamples to III, since the pronouns in the antecedent and elided VPs in each case would be referential and would refer to different people.

The trouble is, however, that there has also been at least one analysis of these examples that is consistent with III: Tomioka (1999) and Elbourne (2008b) have analysed the pronouns in question as identical definite descriptions in each case, basing their analyses on the theories that see donkey pronouns (and related devices) as definite descriptions (Cooper 1979, Heim 1990, Neale 1990, Elbourne 2005). For example, Tomioka (1999) would see the overt pronoun *him* in (11) as meaning something like 'the person x arrested', with the 'x' (informally speaking) bound by the subject *the policeman who arrested John*; we then understand exactly the same meaning for the pronoun in the ellipsis site, with the 'x' being rebound, as it were, by the new subject *the policeman who arrested Bill*. Thus the right meaning would be obtained by means of identical content in antecedent and elided VP. (12) would work analogously, with a meaning for the pronoun something like 'the person having trouble at school'. Since the antecedent and elided VPs in these examples would have identical content, they would not constitute counterexamples to III.

It is beyond the scope of this article to argue in favour of one or another of these analyses of (11) and (12). For current purposes, we can just note that there is at least one plausible analysis of these examples that makes them entirely consistent with III; they are not good sources for counterexamples to this principle, then.

The following example (Jacobson 2018: p. 573) might seem to some to be immune to the analysis of sloppy readings that we just looked at, since it involves the pure indexical *me*:

(13)   A: I want Sally to write to me.
        B: I do too.

In this example, A expresses the desire that Sally write to A; and B expresses the desire that Sally write to B. It might appear to be a counterexample to III, then; and indeed it is possible to imagine an analysis of it in the style of Fiengo & May 1994. Furthermore, as I said, the status of *me* as a pure indexical might lead some to be resistant to the idea that the crucial pronoun here could be bound, which was essential to the widely adopted analysis of sloppy readings suggested by Keenan (1971). But this would be to move too fast. It has been known for a long time now, in fact, that first- and second-

person pronouns do sometimes seem to be bound. Partee (1989: footnote 3), for example, gave the following example:

(14)   I'm the only one around here who will admit that I could be wrong.

This displays what appears to be a bound reading of the second occurrence of *I*, whereby the speaker says that he or she is the only person who has the property [$\lambda x.x$ will admit that x could be wrong]. Analyses of this phenomenon have varied (and will be discussed more in the next section), but there is no denying that we have a bound variable here: the only bone of contention is whether we have a genuine occurrence of the word *I* being bound or whether, on the other hand, we have a so-called 'fake indexical', a bindable individual variable that acquires uninterpretable first-person features by syntactic agreement with a higher element (Kratzer 2009). Be that as it may, the availability of bound variable analyses for first- and second-person pronouns opens up the way for claiming that (13) too involves a sloppy reading brought about by binding; the relevant verb phrase would mean something like [$\lambda x.x$ wants Sally to write to x]. This example too, then, is liable to be analysed in ways that involve antecedent and ellipsis site having identical semantic content, which means that it is not a plausible counterexample to III.

I have not come across any examples that are as convincing as (5) and (6) when it comes to finding counterexamples to III. By the very nature of my project, however, I am open to the idea that such examples may one day come to light.

3.2   Other uses of *current*

Two reviewers for *Semantics & Pragmatics* have pointed out that the semantics of *current* is more flexible than we might assume if we just look at the examples in Section 2. Those examples were compatible with this word obtaining the time that figures as part of its semantic contribution just by looking to a contextual parameter — the time of utterance. But the time contributed can also, it seems, be a contextually salient time that is distinct from the time of utterance. One reviewer points to sentences like the following (obtained by means of an Internet search for naturally occurring examples):

(15)   Back then the current government was the British Crown.

And another reviewer points out that *current* can be bound:

(16)   Whenever I arrive in a new country, I always buy a current newspaper to see what is going on.

Does the existence of this semantic flexibility in *current* undermine the argumentation that has been made so far?

I see two ways in which this could happen. The first possibility, pointed out by an anonymous reviewer, is that a friend of Stanley's argument could claim that III only applies to indexicals that cannot be bound. Thus examples involving *current* would not be counterexamples to III.

However, there are several things wrong with this argument. As the same reviewer pointed out, it would be possible for contextualists to claim that vague predicates like *heap* and *tall* could be bound too. In order to make their case strong, friends of Stanley's argument would have to argue against the possibility of vague predicates being bound, something it is not clear they could do. In fact the same reviewer points to the following piece of evidence that vague predicates can be bound:

(17)   Every time I visit a remote population, its leader stands out as being tall.

There is a prominent reading of this example whereby the tallness is relative to the standards of the visited population: when the speaker visited some African pygmies (average height for males < 4 feet 11 inches), their leader was 5 feet 4 inches tall; and when the speaker visited the Dinka (average height for males > 5 feet 11 inches), their leader was 6 feet 4 inches tall; and so on. Given the structure of (17), this presumably means that we are dealing with donkey anaphora, with *tall* interpreted as something like 'tall by the standards of that population'. So it is unlikely that friends of Stanley's argument could argue that vague predicates cannot be bound.

Another problem with this argument is that it is ad hoc and theoretically unmotivated. Why should there be a difference in behaviour between indexicals that can and indexicals that cannot be bound, when the behaviour in question concerns not bound uses but referential uses?

The final problem with this argument that I will consider is that it is not at all clear that there *are* any indexicals that cannot be bound. What candidates are there? Demonstratives can be bound (Elbourne 2008a):

(18)   Donkeys beat easily. In fact, every man who owns a donkey beats that and nothing else.

(19)   Every man who kept a donkey in this field *(gesture at Field A)* and a donkey in that field *(gesture at Field B)* found that this donkey *(gesture at Field A)* did better than that donkey *(gesture at Field B).*

So can *then* and *there*:

(20)   Everyone who held a party on a weekday found that few people wanted to attend then.

(21)   Every woman who arranged to meet a friend of hers in a graveyard found that her friend did not want to go there.

So can *local* and *enemy* and sundry others (Partee 1989, examples amalgamated):

(22)   Every sports fan went to a local bar and faced an enemy and stole a car and abandoned it two hours later fifty miles away and, in a sudden schizophrenic break, began to think that he was someone else (from another town and with a different personality).

So can whatever it is about weather predicates that contributes a location:

(23)   *(The sports fans in question are in widely separated locations.)* Every sports fan also noticed that it was raining; but he had too much on his plate already to be worried about that.

The first- and second-person singular pronouns require a bit more comment. We have already seen examples like the following, which provide *prima facie* evidence that these items can be bound:

(24)   Only I know why I'm crying.

(25)   Only you know what you have to do.

We now just have to note that, just as there are analyses of this phenomenon like Kratzer's (2009), which claim that we are not dealing with real occurrences of these pronouns here, so there are also analyses that claim that we are (Spathas 2009, Jacobson 2012, 2018)—it is just that (for example) the examples in question involve focus values and the person features of pure indexicals do not make it into focus values. At this point, there is no con-

sensus on this issue. Meanwhile, first-person plural pronouns can even be bound *into* (Partee 1989):

(26)   John often comes over for Sunday brunch. Whenever someone else comes over too, we end up playing trios.

Even *today*, *tomorrow*, *yesterday*, *here*, and *now* might be able to be bound, depending on one's analysis of certain classic examples from Nunberg 1993:

(27)   *(It is the Friday before classes begin at a big university.)* The bookstore crowds usually abate about a week from now.

(28)   *(As above.)* Tomorrow is always the biggest party night of the year.

Suppose, however, that some, presumably small, class of indexicals were found that could not be bound. The proposal would then be to limit III to those words. But what could III possibly tell us about indexicals in general when it is limited to a small sub-class of them? What could it tell us about vague predicates on the assumption that they are indexicals?

Let us move on to the second way in which the semantic flexibility of *current* might be thought to undermine the argumentation presented in this article against Stanley's (2003) argument against contextualism. This is that it might be possible to maintain that examples (5) and (6) involve *current* being bound. If this were the case, of course, then these examples would not be counterexamples to III; they would be compatible with III, in just the way that examples (8a) and (13) in Section 3.1 are arguably compatible with III if the sloppy identity they display is analysed in terms of crucial indexicals being bound.

Here, for concreteness, is one way in which this could come about. Suppose there were a binder of time intervals in a high position in the clause that could make sentences characteristic functions of sets of time intervals. Such a thing could possibly bind *current*, assuming that this word takes a time-interval variable as an argument or is otherwise suitably parameterized. (The exact compositional details do not matter here; I assume that they could be worked out in order to be as generous as reasonably possible to this position opposed to mine.) So for (29a), the first conjunct of (5), we might have a denotation something like (29b); and for (30a), the second conjunct of (5), we would have a denotation something like (30b).

(29)   a.   Sarah likes the current government.
       b. $\lambda t.$ Sarah likes at t the government current at t

(30)   a.   Mary likes the current government.
       b. $\lambda t.$ Mary likes at t the government current at t

Then, I assume, the objection would have to claim that these two conjunct de-
notations could be applied to separate topic times, rather as people working
in situation semantics claim that sentence denotations construed as char-
acteristic functions of sets of situations can be applied to topic situations
(Austin 1950, Barwise & Perry 1983). By this means, it would be possible to
obtain the attested result without having *current* refer to one time in the an-
tecedent VP and another one in the elided VP; the semantic contents of the
two VPs would plausibly be identical, just as in the cases of sloppy identity
analysed in Section 3.1.

   I do not know whether all this is possible. But fortunately, I think that no
lengthy theoretical investigation is necessary, since the objection in question
can be disarmed by examining one additional example. As a minimal variant
on (5) (repeated with extraneous material removed as (31a)), consider (31b):

(31)   a.   Sarah likes the current government and…*(speaking very slowly,
            while a sudden coup takes place)* Mary does now.
       b.   Sarah likes the current government now…*(speaking very slowly,
            while a sudden coup takes place)* but doesn't now.

Example (31b), which is as natural as (31a) in the relevant reading, involves not
sentences but some subsentential constituents being conjoined. If we con-
strue sentences as Tense Phrases (TPs), the constituents in question would
presumably be of category T′. Be that as it may, the conjoined constituents in
(31b) are not propositional; that is, they do not express anything capable of
being evaluated for truth or falsity at a time interval or other circumstance
of evaluation. Given this, it does not make sense to claim that the conjuncts
in (31b) could be evaluated with respect to different topic times, because it
does not make sense to claim that these constituents are evaluated with re-
spect to any topic times at all. Distinct topic times were needed, however, by
the current objection in order to have Sarah and Mary liking governments in
power at different times in (31a) without having the word *current* referring
to different times. So the current objection does not succeed.

4   Conclusion

Stanley's (2003) objection to contextualist theories of vagueness relied on the truth of what I called III, since it claimed that the contextualist account of the sorites was inconsistent with this generalization. But III is false. So Stanley's objection to contextualist theories of vagueness is unsound.

There have been other objections to contextualist theories of vagueness, of course.[7] For example, Sorensen (1998) and Keefe (2007) allege that these theories, with their unconscious context-induced changes in semantic content, leave the door open to simple valid-seeming arguments actually being invalid due to undetected equivocation on vague terms. But there is some reason to believe that Stanley's has often been regarded as the most problematic objection to these theories: for instance, as mentioned in footnote 1, it is the only criticism of this kind of contextualism described in detail in the survey article on the sorites by Sainsbury & Williamson (2017). Examining the other objections to contextualist theories of vagueness is beyond the scope of the current article. But contextualists have reason for optimism.

References

Åkerman, Jonas. 2012. Contextualist theories of vagueness. *Philosophy Compass* 7(7). 470–80. https://doi.org/10.1111/j.1747-9991.2012.00495.x.

Austin, J.L. 1950. Truth. *Supplementary Proceedings of the Aristotelian Society* 24. 111–29. https://www.aristoteliansociety.org.uk/pdf/austin.pdf.

Barwise, Jon & John Perry. 1983. *Situations and attitudes*. Cambridge, Mass.: MIT Press.

Cooper, Robin. 1979. The interpretation of pronouns. In Frank Heny & Helmut Schnelle (eds.), *Syntax and semantics 10: Selections from the third Gröningen round table*, 61–92. New York: Academic Press.

Elbourne, Paul. 2005. *Situations and individuals*. Cambridge, Mass.: MIT Press.

Elbourne, Paul. 2008a. Demonstratives as individual concepts. *Linguistics and Philosophy* 31. 409–66. https://doi.org/10.1007/s10988-008-9043-0.

Elbourne, Paul. 2008b. Ellipsis sites as definite descriptions. *Linguistic Inquiry* 39(2). 191–220. https://doi.org/10.1162/ling.2008.39.2.191.

Ellis, Jonathan. 2004. Context, indexicals and the sorites. *Analysis* 64(4). 362–64. https://doi.org/10.1093/analys/64.4.362.

---

7 See Åkerman 2012 and Sorensen 2018 for useful overviews.

Fara, Delia Graff. 2000. Shifting sands: An interest-relative theory of vagueness. *Philosophical Topics* 28. Originally published under the name 'Delia Graff'., 45–81. https://doi.org/10.5840/philtopics20002816.

Fiengo, Robert & Robert May. 1994. *Indices and identity*. Cambridge, Mass.: MIT Press.

Gert, Joshua. 2008. Vague terms, indexicals, and vague indexicals. *Philosophical Studies* 140(3). 437–45. https://doi.org/10.1007/s11098-007-9154-4.

Hardt, Dan. 1999. Dynamic interpretation of verb phrase ellipsis. *Linguistics and Philosophy* 22. 187–221. https://doi.org/10.1023/A:1005427813846.

Heim, Irene. 1990. E-type pronouns and donkey anaphora. *Linguistics and Philosophy* 13. 137–77. https://doi.org/10.1007/BF00630732.

Jacobson, Pauline. 2012. Direct compositionality and 'uninterpretability': The case of (sometimes) uninterpretable features on pronouns. *Journal of Semantics* 29. 305–43. https://doi.org/10.1093/jos/ffs005.

Jacobson, Pauline. 2018. Some people think there is neg raising, and some don't: Neg raising meets ellipsis. *Linguistic Inquiry* 49(3). 559–76. https://doi.org/10.1162/ling_a_00282.

Kamp, Hans. 1981. The paradox of the heap. In Uwe Mönnich (ed.), *Aspects of philosophical logic*, 225–77. Dordrecht: Reidel.

Kaplan, David. 1989. Demonstratives: An essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals. In Joseph Almog, John Perry & Howard Wettstein (eds.), *Themes from Kaplan*, 481–563. Oxford: Oxford University Press.

Keefe, Rosanna. 2007. Vagueness without context change. *Mind* 116. 275–92. https://doi.org/10.1093/mind/fzm275.

Keenan, Edward. 1971. Names, quantifiers, and a solution to the sloppy identity problem. *Papers in Linguistics* 4. 211–32. https://doi.org/10.1080/08351817109370257.

Kratzer, Angelika. 2009. Making a pronoun: Fake indexicals as windows into the properties of pronouns. *Linguistic Inquiry* 40(2). 187–237. https://doi.org/10.1162/ling.2009.40.2.187.

Neale, Stephen. 1990. *Descriptions*. Cambridge, Mass.: MIT Press.

Nunberg, Geoffrey. 1993. Indexicality and deixis. *Linguistics and Philosophy* 16. 1–43. https://doi.org/10.1007/BF00984721.

Partee, Barbara. 1989. Binding implicit variables in quantified contexts. In Caroline Wiltshire, Bradley Music & Randolph Graczyk (eds.), *Papers from CLS 25*, 342–65. Chicago: Chicago Linguistics Society.

Raffman, Diana. 1994. Vagueness without paradox. *Philosophical Review* 103(1). 41–74. https://doi.org/10.2307/2185872.

Raffman, Diana. 2005. How to understand contextualism about vagueness: Reply to Stanley. *Analysis* 65(3). 244–8. https://doi.org/10.1093/analys/65.3.244.

Ross, John. 1967. *Constraints on variables in syntax*. Cambridge, Mass.: MIT dissertation.

Sainsbury, Mark & Timothy Williamson. 2017. Sorites. In Bob Hale, Crispin Wright & Alexander Miller (eds.), *A companion to the philosophy of language*, 2nd, vol. 2, chap. 28, 734–64. Chichester: Wiley. https://doi.org/10.1002/9781118972090.ch28.

Shapiro, Stewart. 2006. *Vagueness in context*. Oxford: Oxford University Press.

Soames, Scott. 1999. *Understanding truth*. Oxford: Oxford University Press.

Soames, Scott. 2002. Replies. *Philosophy and Phenomenological Research* 62. 429–52. https://doi.org/10.1111/j.1933-1592.2002.tb00215.x.

Sorensen, Roy. 1998. Ambiguity, discretion and the sorites. *The Monist* 81. 215–32. https://doi.org/10.5840/monist199881216.

Sorensen, Roy. 2018. Vagueness. In Edward N. Zalta (ed.), *The Stanford encyclopedia of philosophy*, Summer 2018. Metaphysics Research Lab, Stanford University. http://plato.stanford.edu/archives/sum2018/entries/vagueness/.

Spathas, George. 2009. Interpreting gender features on bound pronouns. In Claire Halpert, Jeremy Hartman & David Hill (eds.), *Proceedings of the 2007 workshop in Greek syntax and semantics*. Cambridge, Mass.: MIT Working Papers in Linguistics.

Stanley, Jason. 2003. Context, interest relativity and the sorites. *Analysis* 63(4). 269–80. https://doi.org/10.1093/analys/63.4.269.

Tomioka, Satoshi. 1999. A sloppy identity puzzle. *Natural Language Semantics* 7. 217–41. https://doi.org/10.1023/A:1008309217917.

Wescoat, Michael. 1989. Sloppy readings with embedded antecedents. Unpublished manuscript, Stanford University.

Paul Elbourne
Magdalen College
Oxford OX1 4AU
United Kingdom
paul.elbourne@magd.ox.ac.uk