

Chapter 18

Sorting out the Sorites

David Ripley

18.1 Introduction

Supervaluational theories of vagueness have achieved considerable popularity in the past decades, as seen in e.g., [Fine \(1975\)](#) and [Lewis \(1970\)](#). This popularity is only natural; supervaluations let us retain much of the power and simplicity of classical logic, while avoiding the commitment to strict bivalence that strikes many as implausible.

Like many non-classical logics, the supervaluationist system SP has a natural dual, the subvaluationist system SB, explored in e.g., [Hyde \(1997\)](#) and [Varzi \(2000\)](#).¹ As is usual for such dual systems, the classical features of SP (typically viewed as benefits) appear in SB in ‘mirror-image’ form, and the non-classical features of SP (typically viewed as costs) also appear in SB in ‘mirror-image’ form. Given this circumstance, it can be difficult to decide which of the two dual systems is better suited for an approach to vagueness.²

The present paper starts from a consideration of these two approaches—the supervaluational and the subvaluational—and argues that neither of them is well-positioned to give a sensible logic for vague language. The first section presents the

¹Although there are many different ways of presenting a supervaluational system, I’ll ignore these distinctions here; my remarks should be general enough to apply to them all, or at least all that adopt the so-called ‘global’ account of consequence. (For discussion, see [Varzi 2007](#).) Similarly for subvaluational systems.

²The situation is similar for approaches to the Liar paradox; for discussion, see e.g., [Beall and Ripley \(2004\)](#) and [Parsons \(1984\)](#).

D. Ripley (✉)

Department of Philosophy, University of Connecticut, Storrs, CT, USA

School of Historical and Philosophical Studies, The University of Melbourne, Parkville, VIC, Australia

e-mail: davewripley@gmail.com

systems SP and SB and argues against their usefulness. Even if we suppose that the general picture of vague language they are often taken to embody is accurate, we ought not arrive at systems like SP and SB. Instead, such a picture should lead us to truth-functional systems like strong Kleene logic (K_3) or its dual LP. The second section presents these systems, and argues that supervaluationist and subvaluationist understandings of language are better captured there; in particular, that a dialethic approach to vagueness based on the logic LP is a more sensible approach. The next section goes on to consider the phenomenon of higher-order vagueness within an LP-based approach, and the last section closes with a consideration of the sorites argument itself.

18.2 S'valuations

Subvaluationists and supervaluationists offer identical pictures about how vague language works; they differ solely in their theory of truth. Because their overall theories are so similar, this section will often ignore the distinction between the two; when that's happening, I'll refer to them all as s'valuationists. First I present the shared portion of the s'valuationist view, then go on to lay out the difference between subvaluational and supervaluational theories of truth, and offer some criticism of both the subvaluationist and supervaluationist approaches.

18.2.1 *The Shared Picture*

It's difficult to suppose that there really is a single last noonish second, or a single oldest child, & c.³ Nonetheless, a classical picture of negation seems to commit us to just that. After all, for each second, the law of excluded middle, $A \vee \neg A$, tells us it's either noonish or not noonish, and the law of non-contradiction, $\neg(A \wedge \neg A)$, tells us it's not both. Now, let's start at noon (since noon is clearly noonish) and move forward second by second. For a time, the seconds are all noonish, but the classical picture seems to commit us to there being a second—just one second!—that tips the scale over to non-noonishness.

Many have found it implausible to think that our language is pinned down that precisely, and some of those who have found this implausible (e.g., [Fine 1975](#)) have found refuge in an s'valuational picture. The key idea is this: we keep that sharp borderline (classicality, as noted above, seems to require it), but we allow that there are many different places it might be. The s'valuationists then take vagueness to be something like ambiguity: there are many precise extensions that a vague

³It's not un-supposable, though; see e.g., [Sorensen \(2001\)](#) and [Williamson \(1994\)](#) for able defences of such a position.

predicate might have, and (in some sense) it has all of them.⁴ It's important that this 'might' not be interpreted epistemically; the idea is not that one of these extensions is the one, and we just don't know which one it is. Rather, the idea is that each potential extension is part of the meaning of the vague predicate. Call these potential extensions 'admissible precisifications'.

The phenomenon of vagueness involves a three-part structure *somehow*; on this just about all theorists are agreed. Examples: for a vague predicate F , epistemicists (e.g., Williamson 1994) consider (1) things that are known to be F , (2) things that are known not to be F , and (3) things not known either to be F or not to be F ; while standard fuzzy theorists (e.g., Machina 1972; Smith 2008) consider (1) things that are absolutely F , or F to degree 1, (2) things that are absolutely not F , or F to degree 0, and (3) things that are neither absolutely F nor absolutely not F . S'valuationists also acknowledge this three-part structure: they talk of (1) things that are F on every admissible precisification, (2) things that are not- F on every admissible precisification, and (3) things that are F on some admissible precisifications and not- F on others.

Consider 'noonish'. It has many different precisifications, but there are some precisifications that are admissible and others that aren't. (37a)–(38b) give four sample precisifications; (37a) and (37b) are admissible precisifications for 'noonish', but (38a) and (38b) aren't:⁵

- (37) a. $\{x : x \text{ is between } 11:40 \text{ and } 12:20\}$
 b. $\{x : x \text{ is between } 11:45 \text{ and } 12:30\}$
- (38) a. $\{x : x \text{ is between } 4:00 \text{ and } 4:30\}$
 b. $\{x : x \text{ is between } 11:40 \text{ and } 11:44, \text{ or } x \text{ is between } 12:06 \text{ and } 12:10, \text{ or } x \text{ is between } 12:17 \text{ and } 12:22\}$

There are at least a couple ways, then, for a precisification to go wrong, to be inadmissible. Like (38a), it might simply be too far from where the vague range is; or like (38b), it might fail to respect what are called *penumbral connections*. 12:22 might not be in every admissible precisification of 'noonish', but if it's in a

⁴NB: S'valuationists differ in the extent to which they take vagueness to be like ambiguity, but they all take it to be like ambiguity in at least this minimal sense. Smith (2008) draws a helpful distinction between supervaluationism and what Smith calls 'plurivaluationism'. Although both of these views have both travelled under the name 'supervaluationism', they are distinct. Supervaluationism makes use of non-bivalent semantic machinery (for example, the machinery in Fine 1975), while plurivaluationism makes do with purely classical models, insisting merely that more than one of these models is the intended one. My discussion of supervaluationism here is restricted to the view Smith calls supervaluationism.

⁵At least in most normal contexts. Vague predicates seem particularly context-sensitive, although they are not the only predicates that have been claimed to be (see e.g., Recanati 2004; Wilson and Sperber 2002). For the purposes of this paper, I'll assume a single fixed (non-wacky) context; these are theories about what happens *within* that context. Some philosophers (e.g., Raffman 1994) have held that taking proper account of context is itself sufficient to dissolve the problems around vagueness. I disagree, but won't address the issue here.

certain admissible precisification, then 12:13 ought to be in that precisification too. After all, 12:13 is more noonish than 12:22 is. Since this is a connection within the penumbra of a single vague predicate, we can follow [Fine \(1975\)](#) in calling it an *internal* penumbral connection.

Admissible precisifications also must respect *external* penumbral connections. The key idea here is that the extensions of vague predicates sometimes depend on each other. Consider the borderline between green and blue. It's sometimes claimed that something's being green rules out its also being blue. If this is so, then no admissible precisification will count a thing as both green and blue, even if some admissible precisifications count it as green and others count it as blue. In order to handle this phenomenon in general, it's crucial that we count precisifications not as precisifying one predicate at a time, but instead as precisifying multiple predicates simultaneously. That way, they can give us the requisite sensitivity to penumbral connections.

18.2.1.1 S'valuational Models

Now that we've got the core of the idea down, let's see how it can be formally modeled.⁶ An SV model M is a tuple $\langle D, I, P \rangle$ such that D , the domain, is a set of objects; I is a function from terms in the language to members of D ; and P is a set of precisifications: functions from predicates to subsets of D .⁷

We then extend each precisification $p \in P$ to a valuation of the full language. For an atomic sentence Fa , a precisification $p \in P$ assigns Fa the value 1 if $I(a) \in p(F)$, 0 otherwise. This valuation of the atomics is extended to a valuation of the full language (in $\wedge, \vee, \neg, \forall, \exists$) in the familiar classical way. In other words, each precisification is a full classical valuation of the language, using truth values from the set $V_0 = \{1, 0\}$.

Then the model M assigns a value to a sentence simply by collecting into a set the values assigned to the sentence by M 's precisifications: $M(A) = \{v \in V_0 : \exists p \in P(p(A) = v)\}$. Thus, models assign values to sentences from the set $V_1 = \{\{1\}, \{0\}, \{1, 0\}\}$. Note that so far, these values are uninterpreted; they work merely as a record of a sentence's values across precisifications. $M(A) = \{1\}$ iff A gets value 1 on every precisification in M ; $M(A) = \{0\}$ iff A gets value 0 on every precisification in M ; and $M(A) = \{1, 0\}$ iff A gets value 1 on some precisifications in M and 0 on others.

⁶There are many ways to build s'valuational models. In particular, one might not want to have to *fully* precisify the language in order to assign truth-values to just a few sentences. Nonetheless, the approach to be presented here will display the logical behaviour of s'valuational approaches, and it's pretty simple to boot. So we can get the picture from this simple approach.

⁷And from propositional variables directly to classical truth-values, if one wants bare propositional variables in the language. Vague propositional variables can be accommodated in this way as well as precise ones.

18.2.2 Differences in Interpretation and Consequence

The s'-valuationists agree on the interpretation of two of these values: {1} and {0}. If $M(A) = \{1\}$, then A is true on M . If $M(A) = \{0\}$, then A is false on M . But the subvaluationist and the supervaluationist differ over the interpretation of the third value: {1,0}.

The supervaluationist (of the sort I'm interested in here) claims that a sentence must be true on *every* precisification in order to be true *simpliciter*, and false on *every* precisification in order to be false *simpliciter*. Since the value {1,0} records a sentence's taking value 1 on some precisifications and 0 on others, when $M(A) = \{1, 0\}$, A is neither true nor false on M for the supervaluationist.

The subvaluationist, on the other hand, claims that a sentence has only to be true on *some* precisification to be true *simpliciter*, and false on *some* precisification in order to be false *simpliciter*. So, when $M(A) = \{1, 0\}$, the subvaluationist says that A is both true and false on M .

Both define consequence in the usual way (via truth-preservation):

- (39) $\Gamma \vDash \Delta$ iff, for every model M , either δ is true on M for some $\delta \in \Delta$, or γ fails to be true on M for some $\gamma \in \Gamma$.⁸

Since the subvaluationist and the supervaluationist differ over which sentences are true on a given model (at least if the model assigns {1,0} anywhere), this one definition results in two different consequence relations; call them \vDash_{SB} (for the subvaluationist) and \vDash_{SP} (for the supervaluationist).

18.2.2.1 \vDash_{SB} and \vDash_{SP}

One striking (and much-advertised) feature of these consequence relations is their considerable classicality. For example (where \vDash_{CL} is the classical consequence relation):

- (40) $\vDash_{\text{CL}} A$ iff $\vDash_{\text{SB}} A$ iff $\vDash_{\text{SP}} A$

- (41) $A \vDash_{\text{CL}}$ iff $A \vDash_{\text{SB}}$ iff $A \vDash_{\text{SP}}$

(40) tells us that these three logics have all the same logical truths, and (41) tells us that they have all the same explosive sentences.⁹ What's more, for any classically

⁸Note that this is a multiple-conclusion consequence relation. One can recover a single-conclusion consequence relation from this if one is so inclined, but for present purposes the symmetrical treatment will be more revealing. See e.g., Restall (2005) for details, or Hyde (1997) for application to s'-valuations. See also Keefe (2000) for arguments against using multiple-conclusion consequence, and Hyde (2010) for response.

⁹Explosive sentences are sentences from which one can derive any conclusions at all, just as logical truths are sentences that can be derived from any premises at all. It's a bit sticky calling them 'logical falsehoods', as may be tempting, since some sentences (in SB at least) can be false without

valid argument, there is a corresponding SB-valid or SP-valid argument:

$$(42) \quad A_1, \dots, A_i \vDash_{\text{CL}} B_1, \dots, B_j \text{ iff } A_1 \wedge \dots \wedge A_i \vDash_{\text{SB}} B_1, \dots, B_j$$

$$(43) \quad A_1, \dots, A_i \vDash_{\text{CL}} B_1, \dots, B_j \text{ iff } A_1, \dots, A_i \vDash_{\text{SP}} B_1 \vee \dots \vee B_j$$

Let's reason through these a bit.¹⁰ Suppose $A_1, \dots, A_i \vDash_{\text{CL}} B_1, \dots, B_j$. Then, since every precisification is classical, every precisification p (in every model) that verifies all the A s will also verify one of the B s. Consider the same argument subvaluationally; one might have all the premises true in some model (because each is true on some precisification or other), without having all the premises true in the *same* precisification; thus, there's no guarantee that any of the B s will be true on any precisification at all. On the other hand, if one simply conjoins all the premises into one big premise, then if it's true in a model at all it guarantees the truth of all the A s on (at least) a single precisification, and so one of the B s must be true on that precisification, hence true in the model.

Similar reasoning applies in the supervaluational case. If all the A s are true in a model, then they're all true on every precisification; nonetheless it might be that none of the B s is true on every precisification; all the classical validity guarantees is that each precisification has some B or other true on it. But when we disjoin all the B s into one big conclusion, that disjunction must be true on every precisification, so the argument is SP-valid.

Note that (42) and (43) guarantee that \vDash_{SB} matches \vDash_{CL} on single-premise arguments, and that \vDash_{SP} matches \vDash_{CL} on single-conclusion arguments. It is apparent that there is a close relationship between classical logic, subvaluational logic, and supervaluational logic. What's more, for every difference between SB and CL, there is a dual difference between SP and CL, and vice versa. This duality continues as we turn to the logical behaviour of the connectives:

$$(44) \quad \begin{array}{l} \text{a. } A, B \not\vDash_{\text{SB}} A \wedge B \\ \text{b. } A, \neg A \not\vDash_{\text{SB}} A \wedge \neg A \end{array}$$

$$(45) \quad \begin{array}{l} \text{a. } A \vee B \not\vDash_{\text{SP}} A, B \\ \text{b. } A \vee \neg A \not\vDash_{\text{SP}} A, \neg A \end{array}$$

(44a) and (44b) are dual to (45a) and (45b). It is often remarked about supervaluations that it's odd to have a disjunction be true when neither of its disjuncts is, but this oddity can't be expressed via \vDash_{SP} in a single-conclusion format.¹¹ Here, in

failing to be true. And I want to shy away from 'contradiction' here too, since I understand by that a sentence of the form $A \wedge \neg A$, and such a sentence will be explosive here but not in the eventual target system.

¹⁰Here I prove only the LTR directions, but both directions indeed hold; see [Hyde \(1997\)](#) for details.

¹¹This, essentially, is Tappenden's 'objection from upper-case letters' ([Tappenden 1993](#)). With multiple conclusions, there's no need for upper-case letters; the point can be made in any typeface you like.

a multiple-conclusion format, it becomes apparent that this oddity is an oddity in the supervaluational consequence relation, not just in its semantics. And of course, there is a parallel oddity involving conjunction for the subvaluationist.

Disjunction and conjunction can be seen as underwriting existential and universal quantification, respectively, so it is no surprise that the oddities continue when it comes to quantification. A sample:

(46) a. $Fa, Fb, \forall x(x = a \vee x = b) \not\vdash_{\text{SB}} \forall x(Fx)$

(47) a. $\exists x(Fx), \forall x(x = a \vee x = b) \not\vdash_{\text{SP}} Fa, Fb$

The cause is as before: in the SB case there is no guarantee that the premises are true on the same precisification, so they cannot interact to generate the conclusion; while in the SP case there is no guarantee that the same one of the conclusions is true on every precisification, so it may be that neither is true *simpliciter*. In the supervaluational case, consequences for quantification have often been noted;¹² but of course they have their duals for the subvaluationist.

18.2.2.2 What to Say About Borderline Cases?

This formal picture gives rise to certain commitments about borderline cases. (I assume here, and throughout, that every theorist is committed to all and only those sentences they take to be true.) Assume that 12:23 is a borderline case of ‘noonish’. The subvaluationist and the supervaluationist agree in their acceptance of (48a)–(48b), and in their rejection of (49a)–(49b):

(48) a. 12:23 is either noonish or not noonish.

b. It’s not the case that 12:23 is both noonish and not noonish.

(49) a. It’s not the case that 12:23 is either noonish or not noonish.

b. 12:23 is both noonish and not noonish.

On the other hand, they disagree about such sentences as (50a)–(50b):

(50) a. 12:23 is noonish.

b. 12:23 is not noonish.

The subvaluationist accepts both of these sentences, despite her rejection of their conjunction, (49b). On the other hand, the supervaluationist rejects them both, despite her acceptance of their disjunction, (48a). So the odd behaviour of conjunction and disjunction observed above isn’t simply a theoretical possibility; these connectives misbehave every time there’s a borderline case of any vague predicate.

¹²For example, consider the sentence ‘There is a last noonish second’. It is true for the supervaluationist, but there is no second x such that ‘ x is the last noonish second’ is true for the subvaluationist.

A major challenge for either the subvaluationist or the supervaluationist is to justify their deviant consequence relations, especially their behaviour around conjunction and universal quantification (for the subvaluationist) or disjunction and existential quantification (for the supervaluationist). At least *prima facie*, one would think that a conjunction is true iff both its conjuncts are, or that a disjunction is true iff one disjunct is, but the s'valuationists must claim that these appearances are deceiving.

The trouble is generated by the lack of truth-functional conjunction and disjunction in these frameworks. Consider the subvaluational case. If A is true, and B is true, we'd like to be able to say that $A \wedge B$ is true. In some cases we can, but in other cases we can't. The value of $A \wedge B$ depends upon more than just the value of A and the value of B ; it also matters how those values are related to each other precisification to precisification. It's this extra dependence that allows s'valuational approaches to capture 'penumbral connections', as argued for in [Fine \(1975\)](#). Unfortunately, it gets in the way of sensible conjunctions and disjunctions.

18.3 LP and K_3

This trouble can be fixed as follows: we keep the s'valuationist picture for atomic sentences, but then use familiar truth-functional machinery to assign values to complex sentences. This will help us retain more familiar connectives, and allow us to compute the values of conjunctions and disjunctions without worrying about which particular conjuncts or disjuncts we use.

18.3.1 *The Shared Picture*

The informal picture, then, is as follows: to evaluate atomic sentences, we consider all the ways in which the vague predicates within them can be precisified. For compound sentences, we simply combine the values of atomic sentences in some sensible way. But what sensible way? Remember, we're going to end up with three possible values for our atomic sentences— $\{1\}$, $\{0\}$, and $\{1,0\}$ —so we need sensible three-valued operations to interpret our connectives. Here are some minimal desiderata for conjunction, disjunction, and negation:

(51) Conjunction:

- a. $A \wedge B$ is true iff both A and B are true.
- b. $A \wedge B$ is false iff either A is false or B is false.

(52) Disjunction:

- a. $A \vee B$ is true iff either A is true or B is true.
- b. $A \vee B$ is false iff both A and B are false.

(53) Negation:


- a. $\neg A$ is true iff A is false.
- b. $\neg A$ is false iff A is true.

These desiderata alone rule out the s'-valuationist options: as is pointed out in [Varzi \(2000\)](#), SB violates the RTL directions of (51a) and (52b), while SP violates the LTR directions of (51b) and (52a).

As before, we'll have two options for interpreting $\{1,0\}$: we can take it to be both true and false, like the subvaluationist, or we can take it to be neither true nor false, like the supervaluationist. Since the above desiderata are phrased in terms of truth and falsity, it might seem that we need to settle this question before we find appropriate operations to interpret our connectives. It turns out, however, that the same set of operations on values will satisfy the above desiderata whichever way we interpret $\{1,0\}$.

18.3.1.1 LP/ K_3

These are the operations from either strong Kleene logic (which I'll call K_3 ; see e.g., [Kripke 1975](#)) or Priest's Logic of Paradox (which I'll call LP; see e.g., [Priest 1979](#)). Consider the following lattice of truth values:



Take \wedge to be greatest lower bound, \vee to be least upper bound, and \neg to reverse order (it takes $\{1\}$ to $\{0\}$, $\{0\}$ to $\{1\}$, and $\{1,0\}$ to itself). Note that these operations satisfy (51a)–(53b); this is so whether $\{1,0\}$ is interpreted as both true and false or as neither true nor false. For example, consider (52a). Suppose we interpret $\{1,0\}$ LP-style, as both true and false. Then a disjunction is true (has value $\{1\}$ or $\{1,0\}$) iff one of its disjuncts is: RTL holds because disjunction is an *upper* bound, and LTR holds because disjunction is *least* upper bound. On the other hand, suppose we interpret $\{1,0\}$ K_3 -style, as neither true nor false. Then a disjunction is true (has value $\{1\}$) iff one of its disjuncts is: again, RTL because disjunction is an *upper* bound and LTR because it's *least* upper bound. Similar reasoning establishes all of (51a)–(53b). So the LP/ K_3 connectives meet our desiderata.

18.3.1.2 Differences in Interpretation and Consequence

There are still, then, two approaches being considered. One, the K_3 approach, interprets sentences that take the value $\{1,0\}$ on a model as neither true nor false on that model. The other, the LP approach, interprets these sentences as both true and false on that model. This section will explore the consequences of such a difference.

As before, consequence for both approaches is defined as in (39) (repeated here as (54)):

$$(54) \quad \Gamma \vDash \Delta \text{ iff, for every model } M, \text{ either } \delta \text{ is true on } M \text{ for some } \delta \in \Delta, \text{ or } \gamma \text{ fails to be true on } M \text{ for some } \gamma \in \Gamma.$$

And as before, differences in interpretation of the value $\{1,0\}$ result in differences about ‘true’, and so different consequence relations (written here as \vDash_{K_3} and \vDash_{LP}).

First, we should ensure that the connectives behave appropriately, as indeed they do, in both K_3 and LP:

$$(55) \quad \begin{array}{l} \text{a. } A, B \vDash_{LP} A \wedge B \\ \text{b. } A \vee B \vDash_{LP} A, B \end{array}$$

$$(56) \quad \begin{array}{l} \text{a. } A, B \vDash_{K_3} A \wedge B \\ \text{b. } A \vee B \vDash_{K_3} A, B \end{array}$$

As you’d expect given this, so do universal and existential quantification:

$$(57) \quad \begin{array}{l} \text{a. } Fa, Fb, \forall x(x = a \vee x = b) \vDash_{LP} \forall x(Fx) \\ \text{b. } \exists x(Fx), \forall x(x = a \vee x = b) \vDash_{LP} Fa, Fb \end{array}$$

$$(58) \quad \begin{array}{l} \text{a. } Fa, Fb, \forall x(x = a \vee x = b) \vDash_{K_3} \forall x(Fx) \\ \text{b. } \exists x(Fx), \forall x(x = a \vee x = b) \vDash_{K_3} Fa, Fb \end{array}$$

Both consequence relations have other affinities with classical consequence, although neither is fully classical:

$$(59) \quad \begin{array}{l} \text{a. } \vDash_{CL} A \text{ iff } \vDash_{LP} A \\ \text{b. } A \vDash_{CL} \text{ iff } A \vDash_{K_3} \end{array}$$

$$(60) \quad \begin{array}{l} \text{a. } A, \neg A \not\vDash_{LP} B \\ \text{b. } \neg A, A \vee B \not\vDash_{LP} B \end{array}$$

$$(61) \quad \begin{array}{l} \text{a. } A \not\vDash_{K_3} B, \neg B \\ \text{b. } A \not\vDash_{K_3} A \wedge B, \neg B \end{array}$$

(59a) tells us that LP and classical logic have all the same logical truths, while (59b) tells us that K_3 and classical logic have all the same explosive sentences. (60) shows us some of the non-classical features of \vDash_{LP} ; note that the failure of Explosion in (60a) does not come about in the same way as in SB (by failing adjunction), since adjunction is valid in LP, as recorded in (55a). (60b) points out the much-remarked failure of Disjunctive Syllogism in LP. Dual to these non-classicalities are the non-classicalities of K_3 given in (61).

18.3.2 *Vagueness and Ambiguity*

As we've seen, one clear reason to think that LP and K_3 are better logics of vagueness than SB and SP is the sheer sensibleness of their conjunction and disjunction, which SB and SP lacked. LP and K_3 thus allow us to give an s'-valuation-flavoured picture of vague *predication* that doesn't interfere with a more standard picture of *connectives*. But there's another reason why at least some s'-valuationists should prefer the truth-functional approach recommended here, having to do with ambiguity.

As we've seen, the s'-valuational picture alleges at least some similarities between vagueness and ambiguity: at a bare minimum, they both involve a one-many relation between a word and its potential extensions. Some s'-valuationists (e.g., Keefe 2000) stop there, but others (e.g., Fine 1975 in places, Lewis 1982) go farther, claiming that vagueness is actually a species of ambiguity. For these authors, there is an additional question worth facing: what's ambiguity like?

18.3.2.1 Non-uniform Disambiguation

Here's one key feature of ambiguity: when an ambiguous word occurs twice in the same sentence, it can be disambiguated in different ways across its occurrences. For example, consider the word 'plant', which is ambiguous between (at least) *vegetation* and *factory*. Now, consider the sentence (62):

(62) Jimmy ate a plant, but he didn't eat a plant.

It's clear that (62) has a non-contradictory reading; in fact, it has two, assuming for the moment that 'plant' is only two-ways ambiguous. 'Plant' can take on a different disambiguation at each of its occurrences, even when those occurrences are in the same sentence. If this were not the case, if multiple occurrences of an ambiguous word had to be disambiguated uniformly within a sentence, then the standard method of resolving an apparent contradiction—by finding an ambiguity—couldn't work. But of course this method does work.

Now, suppose we wanted to build formal models for an ambiguous language. They had better take this fact into account. But SB and SP cannot—they precisify whole sentences at once, uniformly. Hence, SB and SP could not work as logics for ambiguous language.¹³

LP and K_3 , on the other hand, do not have this bad result. They deal with each occurrence of an ambiguous predicate (each atomic sentence) separately, and combine them truth-functionally. Thus, they avoid the bad consequences faced by s'-valuational pictures. In fact, it is LP that seems to be a superior logic of

¹³Pace Fine (1975), which, in a footnote, proposes SP as a logic for ambiguous language. As noted above, this would make it impossible to explain how one resolves a contradiction by finding an ambiguity—a very bad result.

ambiguous language. Here's why: typically, for an ambiguous sentence to be true, it's not necessary that *every* disambiguation of it be true; it suffices that *some* disambiguation is.¹⁴

Since it's clear that LP and K_3 (and LP in particular) are better logics of ambiguity than SB and SP, those s'-valuationists who take vagueness to be a species of ambiguity have additional reason to adopt LP and K_3 .

18.3.2.2 Asynchronous Precisification

For an s'-valuationist who does not take vagueness to be a species of ambiguity, the above argument applies little direct pressure to use LP or K_3 , but it raises an interesting issue dividing the truth-functional approaches from the s'-valuational approaches: when multiple vague predicates occur in the same sentence, how do the various precisifications of one interact with the various precisifications of the other?

Take the simplest case, where a single vague predicate occurs twice in one sentence. What are the available precisifications of the whole sentence? Suppose a model with n precisifications. On the s'-valuational pictures, there will be n precisifications for the whole sentence; while on a truth-functional picture there will be n^2 ; every possible combination of precisifications of the predicates is available. This can be seen in the LP/ K_3 connectives; for example, where \wedge_0 is classical conjunction,

$$(63) \quad M(A \wedge B) = \{a \wedge_0 b : a \in M(A), b \in M(B)\}$$

$M(A)$ and $M(B)$, recall, are sets of classical values. $M(A \wedge B)$ is then obtained by pulling a pair of classical values, one from $M(A)$ and one from $M(B)$, conjoining these values, and repeating for every possible combination, then collecting all the results into a set. In other words, on this picture, every precisification 'sees' every other precisification in a compound sentence formed with \wedge ; multiple predicates are not precisified in lockstep. The same holds, *mutatis mutandis*, for \vee .

18.3.3 What to Say About Borderline Cases?

So much for the logical machinery. What do these approaches say about borderline cases of a vague predicate? Suppose again that 12:23 is a borderline case of 'noonish'. Consider the following list of claims:

- (64) a. 12:23 is noonish.
b. 12:23 is not noonish.

¹⁴Lewis (1982) argues for LP, in particular, as a logic of ambiguity, and mentions vagueness as one sort of ambiguity.

- c. 12:23 is both noonish and not noonish.
- d. 12:23 is neither noonish nor not noonish.
- e. It's not the case that 12:23 is both noonish and not noonish.
- f. 12:23 is either noonish or not noonish.

All of these are familiar things to claim about borderline cases, although a common aversion to contradictions among philosophers means that some of them, like (64c), are more likely to be heard outside the classroom than in it.¹⁵ All these claims receive the value $\{1,0\}$ in a model that takes 12:23 to be a borderline case of noonish. The LP partisan, then, will hold all of these to be true, while the K_3 partisan will hold none of them to be true. Which interpretation of $\{1,0\}$ is more plausible, then? If we are to avoid attributing massive error to ordinary speakers (and ourselves, a great deal of the time), the LP story is far superior. Accordingly, for the remainder of the paper I'll focus in on LP, setting K_3 aside (although much of what follows holds for K_3 as well as LP).

18.4 Higher-Order Vagueness

Some objections to any dialethic approach to vagueness are considered and ably answered in [Hyde and Colyvan \(2008\)](#). But one objection not considered there might seem to threaten any three-valued approach to vagueness, and in particular the LP approach I've presented: the phenomenon of higher-order vagueness. This section evaluates the LP approach's response to the phenomenon, focusing first on the case of second-order vagueness, and then generalizing the response to take in higher orders as well.

18.4.1 Second-Order Vagueness

So far, we've seen a plausible semantics for vague predicates that depends crucially on the notion of an 'admissible precisification'. A vague atomic sentence is true iff it's true on some admissible precisification, false iff false on some admissible precisification. But which precisifications are admissible? Consider 'noonish'. Is a precisification that draws the line at 12:01 admissible, or is it too early? It has seemed to many theorists (as it seems to me) that 'admissible' in this use is itself vague. Thus, we've run into something of the form of a revenge paradox: theoretical machinery invoked to solve a puzzle works to solve the puzzle, but then the puzzle reappears at the level of the new theoretical machinery.¹⁶

¹⁵See [Ripley \(2011\)](#) for evidence that ordinary speakers agree with such claims as (64c) and (64d).

¹⁶See e.g., [Beall \(2008\)](#) for a discussion of revenge.

It would be poor form to offer an account of the vagueness of ‘admissible’ that differs from the account offered of the vagueness of ‘noonish’. After all, vagueness is vagueness, and similar problems demand similar solutions.¹⁷ So let’s see how the account offered above applies to this particular case of vagueness.

What do we do when a precisification is borderline admissible—that is, both admissible and not admissible? We consider various precisifications of ‘admissible’. This will kick our models up a level, as it were. Models (call them level-2 models, in distinction from the earlier level-1 models) now determine not sets of precisifications, but sets of *sets of* precisifications. That is, a level-2 model is a tuple $\langle D, I, P_2 \rangle$, where D is again a domain of objects, I is again a function from terms in the language to members of D , and P_2 is a set whose members are sets of precisifications.

Every individual precisification p works as before; it still assigns each atomic sentence A a value $p(A)$ from the set $V_0 = \{1, 0\}$. Every set of precisifications assigns each atomic sentence a value as well: a set P_1 of precisifications assigns to an atomic A the value $P_1(A) = \bigcup_{p \in P_1} \{p(A)\}$. These values come from the set $V_1 = \wp(V_0) - \emptyset = \{\{1\}, \{0\}, \{1, 0\}\}$. That is, sets of precisifications work just like level-1 models, as far as atomics are concerned; they simply collect into a set the values assigned by the individual precisifications. A level-2 model $M = \langle D, I, P_2 \rangle$ assigns to every atomic A a value $M(A) = \bigcup_{P_1 \in P_2} \{P_1(A)\}$. It simply collects into a set the values assigned to A by the sets of precisifications in P_2 , so it assigns values from the seven-membered set $V_2 = \wp(V_1) - \emptyset = \{\{\{1\}\}, \{\{0\}\}, \{\{1, 0\}\}, \{\{1\}, \{0\}\}, \{\{1\}, \{1, 0\}\}, \{\{0\}, \{1, 0\}\}, \{\{1\}, \{0\}, \{1, 0\}\}\}$.

In applications to vagueness, presumably only five of these values will be needed. Not much hangs on this fact in itself, but it will better show the machinery of the theory if we take a moment to see why it’s likely to be so. Let’s look at how level-2 models are to be interpreted. Take a model $M = \langle D, I, P_2 \rangle$. Each member P_1 of P_2 is an admissible precisification of ‘admissible precisification’. Some precisifications, those that are in any admissible precisification of ‘admissible precisification’, will be in every such P_1 . Others, those that are in no admissible precisification of ‘admissible precisification’, will be in no such P_1 . And still others, those precisifications on the borderline of ‘admissible precisification’, will be in some of the P_1 s but not others.

Now let’s turn to ‘noonish’. 12:00 is in every admissible precisification of ‘noonish’, no matter how one precisifies ‘admissible precisification’; 12:23 is in some admissible precisifications but not others, no matter how one precisifies ‘admissible precisification’¹⁸; and 20:00 is in no admissible precisifications of ‘noonish’, no matter how one precisifies ‘admissible precisification’. So far so good—and so far, it all could have been captured with a level-1 model.

¹⁷This is sometimes called the ‘principle of uniform solution’. For discussion, see e.g., Colyvan (2008) and Priest (2002).

¹⁸Again, assume a context where this is true.

But there is more structure to map. Some moment between 12:00 and 12:23—let's say 12:10 for concreteness—is in every admissible precisification of 'noonish' on some admissible precisifications of 'admissible precisification', and in some admissible precisifications of 'noonish' but not others on some admissible precisifications of 'admissible precisification'. And some moment between 12:23 and 20:00—let's say 12:34—is in no admissible precisification of 'noonish' on some admissible precisifications of 'admissible precisification', and in some admissible precisifications of 'noonish' but not others on some admissible precisifications of 'admissible precisification'.

Here's a (very toy)¹⁹ model mapping the above structure:

- (65) a. D = the set of times from 12:00 to 20:00
 b. I = the usual map from time-names to times
 c. $P_2 = \{ \{ \{ 12:00-12:38 \}, \{ 12:00-12:15 \} \}, \{ \{ 12:00-12:25 \}, \{ 12:00-12:08 \} \} \}$

Call this model M . Now let's apply it to some atomic sentences: $M(N12:00) = \{ \{ 1 \} \}$, $M(N12:10) = \{ \{ 1 \}, \{ 1, 0 \} \}$, $M(N12:23) = \{ \{ 1, 0 \} \}$, $M(N12:34) = \{ \{ 1, 0 \}, \{ 0 \} \}$, and $M(N20:00) = \{ \{ 0 \} \}$. One suspects that these five values are all one needs of V_2 for (at least most) vague predicates. In order for a sentence to take the value $\{ \{ 1 \}, \{ 0 \}, \{ 1, 0 \} \}$ on a model, the model would have to be set up so that, depending on the precisification of 'admissible precisification', the sentence could be in all admissible precisifications *or* some but not others *or* none at all. It seems unlikely that many predicates have admissible precisifications that work like this. For a sentence to take the value $\{ \{ 1 \}, \{ 0 \} \}$ on a model, something even weirder would have to happen: the model would have to make it so that, depending on the precisification of 'admissible precisification', the sentence could be either true in all admissible precisifications or false in all of them, but there could be no admissible precisification of 'admissible precisification' that would allow the sentence to be true on some admissible precisifications but not others. This too seems unlikely. So I suspect that only five of the seven members of V_2 are likely to be useful for vague predicates, although (as mentioned above) not much hangs on this.²⁰

There is of course the question of interpretation: which of these values counts as true? Again, we should give the same answer here as in the case of first-order vagueness, to avoid *ad hoc*-ery: a sentence is true iff it's true on some admissible precisification; and so it's true iff it's true on some admissible precisification, for some admissible precisification of 'admissible precisification'. That is, any sentence whose value on a model has a 1 in it anywhere—any sentence whose value isn't $\{ \{ 0 \} \}$ —is true on that model.

¹⁹For simplicity, we look at only one predicate: N for 'noonish'. This set is then a set of sets of precisifications for 'noonish'. Let $\{x-y\}$ be the set of times between x and y inclusive.

²⁰Actually I don't see that anything does.

18.4.1.1 Connectives

So that's how our atomics get their values. Of course, we need some way to assign values to compound sentences as well, and the familiar LP operations (call them \wedge_1 , \vee_1 , and \neg_1) won't work—they're defined only over V_1 , but our atomics take values from V_2 . Fortunately, a simple tweak will work, getting us sensible level-2 operations \wedge_2 , \vee_2 , and \neg_2 defined over V_2 .

Recall one of our earlier observations about the LP connectives: in a conjunction, every precisification of one conjunct sees every precisification of the other conjunct (*mutatis mutandis* for disjunction). We can use this to define our level-2 connectives.

Consider the conjunction of two V_2 values u and v . Remember, values from V_2 are sets of values from V_1 , and we already have well-behaved connectives over V_1 . To get one potential V_1 value for the conjunction, we can pull a V_1 value from u and one from v , and conjoin them. If we do that in every possible way, and collect all the results into a set, we get a V_2 value appropriate to be that value of the conjunction. More formally: $u \wedge_2 v = \{u' \wedge_1 v' : u' \in u, v' \in v\}$. The same idea will work for disjunction— $u \vee_2 v = \{u' \vee_1 v' : u' \in u, v' \in v\}$ —and negation— $\neg_2 u = \{\neg_1 u' : u' \in u\}$. So let's simply adopt these as our level-2 connectives.

18.4.1.2 Consequence

Level-2 models now assign values to every sentence; first the atomics, via the sets of sets of precisifications, and then to all sentences, via the level-2 connectives. What's more, we have a set $D_2 \subseteq V_2$ of *designated* values—values that count as true. (Some of them also count as false, of course.) This means that we're in a position to define level-2 consequence. We do it in the expected way: $\Gamma \models_2 \Delta$ iff, for every level-2 model M , either $M(\delta) \in D_2$ for some $\delta \in \Delta$, or $M(\gamma) \notin D_2$ for some $\gamma \in \Gamma$.

So we have a full logic erected 'up a level' from LP, as it were. At first blush, this might seem like not much of a response to the challenge of first-order vagueness. After all, it seems that we simply abandoned the initial theory and adopted another. That would hardly be a persuasive defence. But in fact that's not quite what's happened; as it turns out, $\models_2 = \models_{LP}$.²¹ We haven't actually abandoned the initial theory—we've just offered an alternative semantics for it, one that fits the structure of second-order vagueness quite naturally. What's more, we haven't had to invoke any special machinery to do it. Simply re-applying the first-order theory to itself yields this result.

18.4.2 Generalizing the Construction

Of course, the above construction only works for second-order vagueness, and there is much more to higher-order vagueness than that. In particular, just as it was vague

²¹For proof, see Priest (1984).

which precisifications are admissible precisifications of ‘noonish’, it’s vague which precisifications are admissible precisifications of ‘admissible precisification’. From the above construction, of course, one can predict the reply: we’ll look at admissible precisifications of ‘admissible precisification’ of ‘admissible precisification’, which is of course itself vague, and so on and so on. Let’s lay out a general picture here.

Let an n -set of precisifications be defined as follows: a 0-set of precisifications is just a precisification, and a $(k + 1)$ -set of precisifications is a set of k -sets of precisifications. Let sets V_n of values be defined as follows: $V_0 = \{1, 0\}$, and $V_{k+1} = \wp(V_k) - \emptyset$. A level- n model M_n is then a tuple $\langle D, I, P_n \rangle$ such that D is a domain of objects, I is a function from terms to members of D , and P_n is an n -set of precisifications. Consider an atomic sentence A . We build up its value $M_n(A)$ as follows: in concert with I , every precisification p assigns A a value $p(A)$ from V_0 , and every $(k + 1)$ -set P_{k+1} of precisifications assigns A a value $P_{k+1}(A) = \bigcup_{P_k \in P_{k+1}} \{P_k(A)\}$ from V_{k+1} . $M_n(A)$ is then just $P_n(A)$. For the level-1 and level-2 cases, this is just the same as the above setup, but of course it extends much farther.

Which values count as true? By parallel reasoning to our earlier cases, any value that contains a 1 at any depth. More precisely, we can define a hierarchy D_n of sets of designated values as follows: $D_0 = \{1\}$, and $D_{k+1} = \{v \in V_{k+1} : \exists u \in v(u \in D_k)\}$.

For the connectives: we define a hierarchy \wedge_n, \vee_n, \neg_n of operations as follows: \wedge_0, \vee_0 , and \neg_0 are simply classical conjunction, disjunction, and negation. For values $u_{k+1}, v_{k+1} \in V_{k+1}$, $u_{k+1} \wedge_{k+1} v_{k+1} = \{u_k \wedge_k v_k : u_k \in u_{k+1}, v_k \in v_{k+1}\}$. That is, the a conjunction of sets of values is the set of conjunctions of values from those sets. Similarly for disjunction: $u_{k+1} \vee_{k+1} v_{k+1} = \{u_k \vee_k v_k : u_k \in u_{k+1}, v_k \in v_{k+1}\}$, and for negation: $\neg_{k+1} u_{k+1} = \{\neg_k u_k : u_k \in u_{k+1}\}$. Again, this gives us just what we had before in the case where $n = 1$ or 2, but extends much farther.

We are now in a position to define a hierarchy of consequence relations \models_n as follows: $\Gamma \models_n \Delta$ iff, for every n -level model M_n , either $M_n(\delta) \in D_n$ for some $\delta \in \Delta$, or $M_n(\gamma) \notin D_n$ for some $\gamma \in \Gamma$. Of course, this simply extends our earlier definition to the new level- n framework.

18.4.3 Good News

Just as in the case of second-order vagueness, this construction allows us to fully map the structure of n th-order vagueness for any n . By collecting up (with a bit of jiggery), we can come to an ω -valued model that fully maps the structure of all higher-order vagueness. What’s more, just as in the second-order case, we haven’t affected our consequence relation at all; for every $n \geq 1$ (including ω), $\models_n = \models_1$.²² This shows us that, although we can fully map this structure, there is in fact no need

²²Proof and details can be found in Priest (1984). Note as well that the result can be iterated past ω into the transfinite; I don’t think that’ll be necessary here, since every new level is created to address the vagueness of some finite predicate.

to for logical purposes; the logic we define remains unchanged. We may as well stick with the simple three-valued version. (Or any other version we like. I like the three-valued version for its simplicity, but if there's some reason to prefer another version, then by all means.) It's worth noting that the defender of K_3 can make precisely the same defence here.

18.5 Conclusion: The Sorites

We should close by examining the familiar paradox that arises from vague language: the sorites. Here's a sample, built on 'noonish' (still written N):

1. $N 12:00$
2. $\forall x[(Nx) \rightarrow (N(x + 0:00:01))]$
3. $N 20:00$

Now, before we do any logic at all, we know a few things about this argument. We know that the first premise is true, and that the conclusion is not. That leaves us just a few options: we can deny the second premise, or we can deny that \rightarrow supports modus ponens.²³

Well, what is \rightarrow , anyway? (This question is raised forcefully in [Beall and Colyvan \(2001\)](#).) There are many possibilities. Each possibility creates a distinct sorites argument. And, while we may think that the sorites argument must be solved in a uniform way no matter which vague predicate it's built on, we certainly should not think that it must be solved in a uniform way no matter which binary connective \rightarrow it's built on.

Some examples. (1) Suppose $A \rightarrow B$ is just $A \wedge B$. Then surely modus ponens is valid for \rightarrow ; after all, B would follow from $A \rightarrow B$ alone on this supposition. But the second premise in the sorites is obviously not true, given this interpretation: it simply claims that every moment, along with the moment one second later, is noonish. That's not at all plausible. So on this interpretation, the sorites is to be answered by rejecting the second premise. (2) Suppose on the other hand that $A \rightarrow B$ is $\neg(A \wedge \neg A)$. Then the second premise is clearly true, at least given most theorists' commitments about the law of non-contradiction (including the commitments of the LP-based approach), but it just as clearly does not support modus ponens. From A and $\neg(A \wedge \neg A)$, we cannot conclude B .

²³Some have accepted the conclusion or rejected the first premise (e.g., [Unger 1979](#)), but to take such a position seriously is to remove much of the sense of 'noonish'. And to take it seriously for every vague predicate would make it very hard indeed to talk truly at all. There are other more radical approaches, too: we might reject transitivity of entailment (as in [Zardini 2008](#) or [Cobrerros et al. 2012](#)), or universal instantiation (as in [Kamp 1981](#)). The LP-based solution offered here keeps to the more conservative side of the street.

Of course, $A \wedge B$ and $\neg(A \wedge \neg A)$ are silly conditionals. But they make a point: whatever our commitment to uniform solution, it does not hold when we vary the key connective in the sorites argument. We are free to reject the second premise for some readings of \rightarrow and deny modus ponens for others, and this does not make our solution non-uniform in any way worth avoiding. We might both reject the premise and deny modus ponens for some readings of \rightarrow , for example if we read $A \rightarrow B$ simply as A . The one thing we cannot do is accept both the premise and the validity of modus ponens, on any single reading of \rightarrow .

LP as presented here includes at best a very weak conditional. Its material conditional, defined as $A \supset B := \neg(A \wedge \neg B)$, does not support modus ponens.²⁴ Given the theory of vague predicates advanced here, the second premise of the sorites is true if we read \rightarrow as \supset . So the present account doesn't run into any trouble on that version of the sorites. What's more, as mentioned in Hyde (2001), the Stoics sometimes used this form of the argument (the form using 'for any moment, it's not the case both that moment is noonish and that one second later isn't noonish'), precisely to avoid debates about the proper analysis of conditionals. If we do the same, no trouble ensues.

On the other hand, the most compelling versions of the sorites use the 'if... then' of natural language. \supset isn't a very promising candidate for an analysis of a natural-language conditional, in LP or out of it, because of the well-known paradoxes of material implication (see e.g., Routley et al. 1982 for details). What is the right analysis of natural-language conditionals is a vexed issue (to say the least!) and not one I'll tackle here, so this is not yet a response to the sorites built on 'if... then'. For now, we can see that the LP-based approach answers the material-conditional version of the sorites handily.

What's more, it embodies the picture of vague language underlying subvaluationist and supervaluationist motivations in a more natural way than SB and SP themselves do. It also verifies much of our ordinary talk about borderline cases, contradictory and otherwise, and provides a satisfying non-*ad hoc* account of higher-order vagueness. In short, LP should be considered a serious contender in the field of non-classical approaches to the phenomenon of vagueness.

Acknowledgements Research partially supported by the French government, Agence Nationale de la Recherche, grant ANR-07-JCJC-0070, program "Cognitive Origins of Vagueness", and by the Spanish government, grant "Borderlineness and Tolerance" ref. FFI2010-16984, MICINN. Many thanks as well to Jc Beall, Rachael Briggs, Mark Colyvan, Dominic Hyde, Joshua Knobe, William Lycan, Ram Neta, Graham Priest, Greg Restall, Keith Simmons, Mandy Simons, and Zach Weber, as well as audiences at the Fourth World Congress of Paraconsistency, the University of Queensland, and Carnegie Mellon University for valuable discussion, insight, and support.

²⁴This is because modus ponens on \supset is equivalent to disjunctive syllogism, which anyone who takes contradictions seriously ought to reject. See Priest (1979) for discussion.

References

- Beall, J.C. 2008. Prolegomenon to future revenge. In *Revenge of the liar: New essays on the paradox*, ed. J.C. Beall, 1–30. Oxford: Oxford University Press.
- Beall, J.C., and M. Colyvan. 2001. Heaps of gluts and hyding the sorites. *Mind* 110: 401–408.
- Beall, J.C., and D. Ripley. 2004. Analetheism and dialetheism. *Analysis* 64(1): 30–35.
- Cobreros, P., P. Égré, D. Ripley, and R. van Rooij. 2012. Tolerant, classical, strict. *Journal of Philosophical Logic* 41(2): 347–385.
- Colyvan, M. 2008. Vagueness and truth. In *From truth to reality: New essays in logic and metaphysics*, ed. H. Duke, 29–40. New York: Routledge.
- Fine, K. 1975. Vagueness, truth, and logic. *Synthese* 30: 265–300.
- Hyde, D. 1997. From heaps and gaps to heaps of gluts. *Mind* 106: 641–660.
- Hyde, D. 2001. A reply to beall and colyvan. *Mind* 110: 409–411.
- Hyde, D. 2010. The prospects of a paraconsistent approach to vagueness. In *Cuts and clouds: Vagueness, its nature, and its logic*, ed. R. Dietz and S. Moruzzi, 385–406. Oxford: Oxford University Press.
- Hyde, D., and M. Colyvan. 2008. Paraconsistent vagueness: Why not? *Australasian Journal of Logic* 6: 107–121.
- Kamp, H. 1981. Two theories about adjectives. In *Aspects of philosophical logic*, ed. U. Mönnich. Dordrecht: Reidel.
- Keefe, R. 2000. *Theories of vagueness*. Cambridge: Cambridge University Press.
- Kripke, S. 1975. Outline of a theory of truth. *Journal of Philosophy* 72(19): 690–716.
- Lewis, D. 1970. General semantics. *Synthese* 22: 18–67.
- Lewis, D. 1982. Logic for equivocators. *Noûs* 16(3): 431–441.
- Machina, K.F. 1972. Vague predicates. *American Philosophical Quarterly* 9: 225–233.
- Parsons, T. 1984. Assertion, denial, and the liar paradox. *Journal of Philosophical Logic* 13: 137–152.
- Priest, G. 1979. Logic of paradox. *Journal of Philosophical Logic* 8: 219–241.
- Priest, G. 1984. Hyper-contradictions. *Logique et Analyse* 107: 237–243.
- Priest, G. 2002. *Beyond the limits of thought*. Oxford: Oxford University Press.
- Raffman, D. 1994. Vagueness without paradox. *The Philosophical Review* 103(1): 41–74.
- Recanati, F. 2004. *Literal meaning*. Cambridge: Cambridge University Press.
- Restall, G. 2005. Multiple conclusions. In *Logic, methodology, and philosophy of science: Proceedings of the twelfth international congress*, ed. P. Hajek, L. Valdes-Villanueva, and D. Westerståhl, 189–205. Kings' College, London.
- Ripley, D. 2011. Contradictions at the borders. In *Vagueness and communication*, ed. R. Nouwen, R. van Rooij, H.-C. Schmitz, and Uli Sauerland, 169–188. Heidelberg: Springer.
- Routley, R., R.K. Meyer, V. Plumwood, and R.T. Brady. 1982. *Relevant logics and their rivals*, vol. 1. Atascadero: Ridgeview.
- Smith, N.J.J. 2008. *Vagueness and degrees of truth*. Oxford: Oxford University Press.
- Sorensen, R. 2001. *Vagueness and contradiction*. Oxford: Oxford University Press.
- Tappenden, J. 1993. The liar and sorites paradoxes: Toward a unified treatment. *The Journal of Philosophy* 90(11): 551–577.
- Unger, P. 1979. There are no ordinary things. *Synthese* 41: 117–154.
- Varzi, A. 2000. Supervaluationism and paraconsistency. In *Frontiers in paraconsistent logic*, ed. D. Batens, C. Mortensen, G. Priest, and J.-P. Van Bendegem, 279–297. Baldock: Research Studies Press.
- Varzi, A. 2007. Supervaluationism and its logics. *Mind* 116: 633–676.
- Williamson, T. 1994. *Vagueness*. London/New York: Routledge.
- Wilson, D., and D. Sperber. 2002. Truthfulness and relevance. *Mind* 111: 583–632.
- Zardini, E. 2008. A model of tolerance. *Studia Logica* 90: 337–368.