DIALETHEISM AND DISTRIBUTED SORITES

BEN BLUMSON

Abstract. Noniterative approaches to the sorites paradox accept single steps of soritical reasoning, but deny that these can be combined into valid chains of soritical reasoning. The distributed sorites is a puzzle designed to undermine noniterative approaches to the sorites paradox, by deriving an inconsistent conclusion using only single steps, but not chains, of soritical reasoning. This paper shows how a dialetheist version of the noniterative approach, the strict-tolerant approach, also solves the distributed sorites paradox, at no further cost, by accepting the inconsistent conclusion.

1. Introduction

Intuitively, plucking one hair from a hirsute man will not make him bald. Giving one dollar to a poor person will not make them rich. And adding one grain of sand to a pile will not it make it a heap. In general, according to the principle of *tolerance*, vague predicates are insensitive to tiny differences.[1] Growing just a millimetre taller, for example, may make someone who is strictly less than two metres tall at least two meters tall, but it won't allow anyone who isn't tall already to become tall. Similarly, adding a dollar to someone's income may take them from one tax threshold to the next, but it won't lift a poor person out of poverty.

Tolerance is intuitive, but leads to paradox. Suppose, for example, that I am not bald. Then according to tolerance, plucking one hair from my head will not make me bald. But then from tolerance again, plucking another hair from my head will not make me bald. But continuing to reason in this way,

Date: September 9, 2023.

[1]See, for example, Égré 2015 and the references therein.

we could show that even plucking all the hairs from my head will not make me bald, which is absurd. Likewise, suppose I am not rich. Then according to repeated applications of tolerance, no matter how many dollars you give me, I will never become rich, which is also absurd. These are instances of the well-known paradox of the heap, also known as the *sorites paradox*.[2]

One response to the sorites paradox, known as *noniteration*, is to accept arguments which use it once or twice on its own, but reject arguments which use it many times over.[3] So, for example, from the fact that 0 grains of sand is not a heap, one can infer that 1 grain of sand is not a heap. And from the fact that 1 grain of sand is not a heap, one can infer that 2 grains of sand is not a heap. In general, from the fact that n grains of sand is not a heap, one can infer that $n + 1$ grains of sand is not a heap. Nevertheless, according to noniteration, one cannot string all these arguments together to show, for example, that $100,000$ grains of sand is not a heap.

The *distributed sorites* paradox is a puzzle designed to undermine noniteration as a response to the original sorites paradox.[4] Here is how the puzzle is presented by Zach Barnett:

> There is a $100,000$ step staircase. The bottom step, Step 1, has one grain of sand. Step 2 has two grains. In general, Step n contains n grains (arranged in a heap where possible). The steps toward the bottom obviously do not contain heaps. The steps toward the top obviously do. With respect to some intermediate steps, it's hard to say.
>
> Now, we tinker with the set up: Remove one grain of sand from each step (except the bottom one), and then add all

[2] See, for example, Hyde and Raffman 2018 for an overview of the sorites paradox.

[3] See, for example, Zardini 2008 and Gaifman 2010. See also Barnett 2019, p. 1073 for further references as well as the name "noniteration".

[4] See Barnett 2019, p. 1073. See also Barnett 2018, pp. 8–9 and Carlson, Jedenheim-Edling, and Johansson 2021 for a closely related puzzle in normative ethics.

> 99,999 of the grains taken to Step 1, and arrange them in a heap (Barnett 2019, p. 1074).

From tolerance, it follows that no heaps were destroyed in the process of tinkering, so including the new heap on Step 1, there is one more heap than there was before. But since every step after tinkering corresponds to a unique step before tinkering with the same number of grains on it, it is also clear that there is not one more heap than there was before. So there both is and is not one more heap, which is a contradiction.

How is the distributed sorites a problem for noniteration in particular? According to Barnett, the difference is that whereas the classical sorites paradox applies tolerance to the same object many times, the distributed sorites paradox applies tolerance only to different objects. As he puts it:

> The driving thought behind No Iteration seemed to be that repeated application of tolerance *to a single object* was problematic – since, in effect, such a procedure would result in a very large change to the object (and no one thinks that tolerance holds for very large changes). But even if tolerance can be restricted so as to block the repeated application to single objects, there seems to be no clear grounds for rejecting its application to many *distinct objects* – so long as tolerance is not applied more than once in each case (Barnett 2019, pp. 1073–4).

Noniteration can avoid the classical sorites paradox by avoiding the application of tolerance to the same object many times. But it cannot avoid the distributed sorites paradox, according to Barnett, since the distributed sorites applies tolerance only to different objects.

But this is not quite right, since even traditional versions of the sorites paradox do apply tolerance to different objects. Consider again, for example, the 100,000 step staircase. The first step does not contain a heap. From tolerance, the second step does not contain a heap. So from tolerance again, the third step does not contain a heap. Continuing to reason in this way, we could show that even the $100,000^{th}$ step does not contain a heap. But even though in this case tolerance is applied only to a different thing at each step, this reasoning is just the same as in the traditional sorites paradox, and proponents of noniteration can reject it on the same grounds.

Instead, the important difference is that the traditional sorites reuses conclusions of previous applications of tolerance as premises in subsequent applications, whereas the distributed sorites does not reuse conclusions of previous applications of tolerance as premises in subsequent applications. In the traditional sorites, one reaches the conclusion that the second step does not contain a heap, and then reuses this conclusion as a premise in an argument that the third step does not contain a heap. But if Barnett is right, then in the distributed sorites, one does not conclude that the second step does not contain a heap, nor reuse this conclusion as a premise in any subsequent argument.

How should proponents of noniteration respond to the distributed sorites paradox? That depends on the details of how the intuition behind noniteration is regimented into a systematic logic. As Haim Gaifman writes "The unsophisticated intuition that bans the stringing of "too many" Sorites conditionals is essentially correct; the trick is to find a smooth, non ad-hoc way of building such constraints into the logic" (Gaifman 2010, p. 12).[5] In this paper I will look at one such already extant regimentation, called *st* for "strict-tolerant", which works by tolerating a less strict kind of truth for conclusions than it does for premises, and thus prevents tolerantly true conclusions from recurring as strictly true premises.[6]

[5]Barnett 2019, p. 1075 concedes that Gaifman's contextualist development of noniteration avoids the distributed sorites paradox, but at the cost of qualifying tolerance.

[6]The strict-tolerant approach was proposed by van Rooij 2011 in response to Zardini 2008, and is elaborated on at length by Cobreros et al. 2012.

Note that the strict-tolerant account was developed primarily as a regimentation of noniteration as a response to the traditional sorites paradox, and is thus a version of the noniterative approach the distributed sorites paradox is intended to undermine. However, I will argue in this paper that insofar as the strict-tolerant approach succeeds as a solution to the traditional sorites paradox, it succeeds just as well as a solution to the distributed sorites paradox. The distributed sorites is a new problem – but in this case it succumbs to an old solution.

As we shall see below, the strict-tolerant approach works by distinguishing three kinds of truth – *strict*, *classical* and *tolerant* – and requiring that the premises of a sound argument be strictly true, whereas the conclusion need only be tolerant true. And as we will also see in section 6 below, according to the strict-tolerant approach, although one or two steps of soritical reasoning is valid, three steps of soritical reasoning is not valid, since one step may fall from strictly to classically true, a second step from classically to tolerantly true, but the third step from tolerantly true to not true at all.

Strict truth differs from classical truth because it admits of "gaps" or, in other words, sentences such that neither they nor their negation are true. So although either 'I am rich' or else 'I am not rich', for example, must be classically true, it may be that neither 'I am rich' nor 'I am not rich' is strictly true. Since the admission of gaps is a familiar idea in treatments of vagueness, it is not too surprising that they are incorporated into the strict-tolerant framework. However, gaps are essential neither to the strict-tolerant approach's solution to the sorites paradox, nor its solution to the distributed sorites paradox.

On the other hand, tolerant truth differs from classical truth because it admits of "gluts" or, in other words, sentences such that both they and their negations are true. So although 'I am rich' and 'I am not rich', for example, cannot both be classically true, it may be that 'I am rich' and 'I am not rich' are both tolerantly true. This commits the strict-tolerant approach to the doctrine of *dialetheism*, according to which some contradictions are true.[7] Although the admission of gluts is a less familiar idea in treatments of vagueness, we will see below that dialetheism is essential to the strict-tolerant solution to both the traditional and distributed sorites.

Since dialetheism may seem like an especially radical departure from classical logic, it's worth pausing to make two points to allay this concern. Firstly, dialetheism about vagueness is motivated not only as a last resort, but also by a symmetry between gaps and gluts, which means that admitting gluts usually turns out to depart from classical logic no more radically than admitting gaps does.[8] Given this motivation, I will not be arguing here that the strict-tolerant approach is better than its rivals, but only that it can handle the distributed sorites paradox just as well as it handles the traditional sorites paradox.

Secondly, it is compatible with everything said here that different logics are appropriate for different purposes.[9] In fact, we will see below that the distinction between strict, classical and tolerant truth naturally gives rise to nine distinct logics, more than one of which may well be considered appropriate for modelling reasoning in the presence of vagueness. Amongst these, I single out the strict-tolerant approach because it provides an interesting regimentation of the intuitions behind tolerance and noniteration. But I do not wish to claim that it is the uniquely best logic for reasoning about vagueness, nor to defend it against any other contenders for that title.

How does combining dialetheism with noniteration resolve the distributed sorites? The answer is surprisingly obvious. The distributed sorites involves

---

[7]See, for example, Priest 2006, p. 4 for this characterisation of dialetheism. For a helpful overview of dialetheism about vagueness see Hyde and Colyvan 2008. See Cobreros et al. 2012, p. 356 for the connection of their approach with dialetheism.

[8]Dominic Hyde, for example, uses this symmetry to argue that supervaluationism is no better than subvaluationism, its dialetheic dual (Hyde 1997; Hyde 2010). Cobreros 2011, on the other hand, argues subvaluationism is better than supervaluationism.

[9]See Beall and Restall 2006, pp. 27–8, for example, on vagueness and logical pluralism.

valid reasoning from true premises to an inconsistent conclusion. Classically in such a situation one must either reject one of the premises or dispute the validity of the reasoning. But it is open to proponents of dialetheism to do neither, and to accept the inconsistent conclusion instead. So whereas in the case of the traditional sorites, dialetheic noniteration rejects the reasoning, in the case of the distributed sorites it accepts the reasoning, but also accepts the inconsistency.

2. The Sorites Paradox

In order to regiment tolerance and the sorites paradox, let us adopt the notation xI_Fy to say that x is indiscriminable from y in respect of F-ness. So if F is 'is tall', then xI_Fy says that x is indiscriminable from y in respect of height. We'll assume throughout that xI_Fy stands for a relation which is reflexive and symmetric, but not necessarily transitive. So in respect of height, Abelard is indiscriminable from Abelard. And if Abelard is indiscriminable from Bob in respect of height, Bob is indiscriminable from Abelard. Nevertheless, even if Abelard's height is indiscriminable from Bob's and Bob's from Cecilia's, Abelard may nevertheless be discriminably taller than Cecilia. We'll also assume that indiscriminability is not itself vague.

Recall that according to tolerance, vague predicates are insensitive to very small or indiscriminable differences. So with this notation in hand, we can formulate the principle of tolerance for a vague predicate F as saying that if x is F and x is indiscriminable from y with respect to F-ness, then y is also F, or formally:

(1) $(\forall x)(\forall y)((Fx \land xI_Fy) \rightarrow Fy)$ Tolerance

So tolerance says that if you're indiscriminable from a tall person in respect of height, for example, then you are yourself tall. Likewise, tolerance says that if I was poor yesterday, and my position today is indiscriminable in respect of wealth to my position yesterday, then I am still poor today.

Given this formulation of tolerance, we can formulate the reasoning in the sorites paradox as follows, with tolerance as an implicit premise:

Definition 1 (Sorites Paradox)**.** *Implicit sorites* are arguments of the form:

(2) Fa_0

(3) $a_0 I_F a_1, a_1 I_F a_2, ..., a_{n-1} I_F a_n$

(4) $\therefore Fa_n$

An *explicit* sorites is an implicit sorites with the addition of thesis 1, tolerance, as an explicit premise.[10]

We distinguish between implicit and explicit sorites since they are treated differently both classically and by the strict-tolerant approach. As we will see in section 5 below, *implicit* sorites are invalid in classical logic, even when we extend it with I_F as additional logical vocabulary. But even without treating I_F as additional logical vocabulary, we can check already that *explicit* sorites arguments are classically valid.

(Notice that the combination of an n-step sorites argument culminating in the intermediate conclusion Fa_n with an m-step sorites argument commencing with the premise Fa_n is effectively an $m+n$-step sorites argument culminating with the conclusion Fa_{m+n}. And in classical logic, of course, if an n-step and an m-step sorites argument are both valid, it follows that the $n+m$ step sorites argument which results from putting these arguments together is valid too.[11] So according to classical logic, if a single-step sorites argument is valid, then sorites arguments of any finite number of steps, no matter how high, are also valid.)

Although explicit sorites are classically valid, in the presence of vagueness it appears possible for sorites arguments to lead from true premises to false conclusions. It may be, for example, that Abelard is tall, that Abelard is

[10] In Cobreros et al. 2012, p. 376, their Version 1 of the sorites is our implicit sorites, and their Version 2 of the sorites is our explicit sorites.

[11] Essentially due to the "cut rule", for which see Ripley 2013a and references therein.

indiscriminable from Bob in respect of height, that Bob is indiscriminable from Cecilia in respect of height, and so on ... until we reach Zach who, because small differences can accumulate to form large ones, is nevertheless short. So in order to accept the validity but reject the soundness of explicit sorites arguments, classical logicians are forced to reject one of the premises, viz.: tolerance.[12] This is the sorites paradox.

How does the strict-tolerant approach address the sorites paradox? As I explain in detail in section 6 below, it will turn out that according to the strict-tolerant approach implicit sorites are valid for $n \leq 2$, but invalid for any $n > 2$, thus regimenting noniteration as a solution to the implicit sorites paradox. However, it will also turn out that the *explicit* sorites is valid for any n whatsoever, and so to solve the explicit sorites paradox, the proponent of the strict-tolerant approach will have to join with the classical logician in rejecting tolerance as a premise of the argument.

However, as I explain in more detail below, because the strict-tolerant approach distinguishes between the kind of truth required for premises and conclusions, they may reject tolerance as a premise of the argument, on the grounds that it is not strictly true, while at the same time accepting it as a tautology, on the grounds that it must be tolerantly true. So arguably, by rejecting tolerance as a premise but accepting it as a tautology, the proponent of the strict-tolerant approach can explain why tolerance is such an intuitive principle, while still joining the classical logician in rejecting the soundness of explicit sorites arguments.

3. The Distributed Sorites

How should we regiment the distributed sorites using the notation of the previous section? Firstly, note that taking a grain from each step to create a new heap is not strictly necessary to generate the puzzle, since it produces no real change in the set-up. As Barnett writes:

[12]See, for example, Sorensen 1988, p. 239.

> The new configuration of sand is – in all relevant respects – the same as the original configuration. We began with $100,000$ distinct heap candidates, ranging in size from 1 to $100,000$. Sure enough, this is precisely what we are left with after moving the grains as prescribed. In effect, we have simply relocated each object: The largest object, which was at the top, is now at the bottom; each other object moved up a step, as if on an escalator (Barnett 2019, p. 1075).

In other words, the device of taking a grain from each step to create a new heap merely serves to dramatise that there is a one-to-one correspondence, or a bijective function, between steps of the staircase such that the first step corresponds to the $100,000^{th}$, the second step corresponds to the first step, ..., the n^{th} step corresponds to the $n-1^{th}$ step, ..., and the $100,000^{th}$ step corresponds to the $99,999^{th}$, so that each step except the first is indiscriminable from the step it corresponds to.

However, if tolerance is true, this correspondence reveals an inconsistency about the number of heaps. Since each step except the first is indiscriminable from the heap it corresponds to, it follows from tolerance that each step which contains a heap corresponds to a step which also contains a heap. But although the first step does not contain a heap, it also corresponds to a step which does contain a heap. So there is a heap corresponding to each heap, plus one more heap corresponding to the first step. In other words, if there are n heaps, then the number of heaps is $n+1$. But this is inconsistent – there cannot be more heaps than there are.

Representing this one-to-one correspondence as a bijective function f, we can regiment the distributed sorites paradox as follows:

Definition 2 (Distributed Sorites Paradox)**.** An *implicit distributed sorites* is an argument of the form:

(5) $(\forall y)(\exists x)f(x) = y \wedge (\forall x)(\forall y)(f(x) = f(y) \rightarrow x = y)$ Bijection

(6) $(\forall x)(x \neq a_0 \rightarrow xI_F f(x))$ Indiscriminability

(7) $(\forall x)(x = a_0 \vee ... \vee x = a_n)$ Finiteness

(8) $Fa_0 \wedge \neg Fa_n \wedge f(a_0) = a_n$ Nontriviality

(9) $\therefore (\exists x)(Fx \wedge \neg Fx)$ Inconsistency

An *explicit* distributed sorites is an implicit distributed sorites with the addition of thesis 1, tolerance, as an explicit premise.

As for the traditional sorites, we distinguish between implicit and explicit distributed sorites because they are treated differently both in classical logic and by the strict-tolerant approach. As we will see in more detail in section 5 below, *implicit* distributed sorites are invalid in classical logic, even when we extend it with I_F as additional logical vocabulary. But even without treating I_F as additional logical vocabulary, we can check already that *explicit* distributed sorites arguments are classically valid. But of course, there is no additional problem for proponents of classical logic here, as they were already committed to denying tolerance for the sake of solving the traditional sorites paradox.

How should the strict-tolerant approach address the distributed sorites paradox? As I explain in detail in section 7 below, it will turn out that according to the strict-tolerant approach the implicit distributed sorites argument is valid for any n whatsoever. But instead of rejecting a premise of the implicit distributed sorites argument, a proponent of the strict-tolerant approach can and should respond by accepting the inconsistent conclusion, since tolerant truth admits of inconsistency. Moreover, as I will emphasise in section 6, the strict-tolerant approach is already committed to this inconsistency in the case of the traditional sorites, and so there is no additional cost to accepting it in the case of the distributed sorites.

However, note that while the proponent of the strict-tolerant approach has a principled reason for accepting the inconsistent conclusion, it does not follow that they may accept any conclusion at all, and so any premises that they accept must be satisfiable or, in other words, admit of interpretations where they are all strictly true. As I will explain in more detail in section 7 below, it turns out that according to the strict-tolerant approach the premises of implicit distributed sorites are unsatisfiable for $n \leq 2$, but are satisfiable for any $n > 2$, and so the strict-tolerant approach can meet this desideratum for resolving the implicit distributed sorites.

But it will also turn out that the premises of *explicit* distributed sorites arguments are not satisfiable for any n, and so to solve the explicit distributed sorites paradox, the proponent of the strict-tolerant approach will again have to reject tolerance as a premise. But as the strict-tolerant approach already had to reject tolerance as a premise to address the traditional sorites paradox, there is no additional cost to rejecting tolerance as a premise in order to address the distributed sorites paradox also. So, I will argue, insofar as the strict-tolerant approach can solve the traditional sorites paradox, it also solves the distributed sorites paradox at no extra cost.

4. Tolerant, Classical, Strict

How proponents of noniteration should respond to the distributed sorites paradox depends not only on how the paradox is regimented, but also on how noniteration itself is incorporated into a comprehensive logical system. In this section I introduce for this purpose two apparently orthogonal ideas about vagueness, viz. that some borderline cases satisfy neither a vague predicate nor its negation or – more importantly for our purposes – that some borderline cases satisfy both a vague predicate and its negation. To make this precise, we introduce a distinction between three kinds of truth for vague statements: *classical*, *tolerant*, and *strict*.

Classical truth, of course, is defined just as it is in classical logic, but then tolerant and strict truth are defined in terms of classical truth and the relation of indiscriminability, via a mutually recursive definition.[13] The base

[13]See van Rooij 2011, p. 213 and Cobreros et al. 2012, p. 353 for this mutually recursive definition of strict and tolerant truth.

clause of this definition covers simple predication: a sentence of the form Fa is tolerantly true if and only if there is b indiscriminable from a with respect to F-ness and Fb is classically true, whereas a sentence of the form Fa is strictly true if and only if for all b indiscriminable from a with respect to F-ness, Fb is classically true.

The recursive clauses for the quantifiers, disjunction and conjunction are as usual, except for negation: a sentence of the form $\neg A$ is tolerantly true if and only if A is not strictly true, whereas $\neg A$ is strictly true if and only if A is not tolerantly true (and as usual, a sentence of the form $\neg A$ is classically true if and only if A is not classically true). In other words, tolerant and strict truth are duals – to be tolerantly true is to be not strictly not true, and to be strictly true is to be not tolerantly not true. Finally, the conditional is defined as material implication, viz.: $A \rightarrow B =_{def} \neg A \vee B$.

Classically, either Fa or $\neg Fa$ must be true. But in borderline cases, it may be that neither Fa nor $\neg Fa$ is strictly true. Suppose, for example, that Fa is classically true, but Fb is not classically true, and $aI_F b$. Then from the base clause for strict truth, Fa is not strictly true. But from the base clause for tolerant truth, Fa is tolerantly true, and so from the recursive clause for negation, $\neg Fa$ is not strictly true. Hence, strict truth admits of gaps, where neither a vague predication nor its negation is true.

Likewise, classically Fa and $\neg Fa$ cannot both be true. But in borderline cases, it may be that both Fa and $\neg Fa$ are tolerantly true. Suppose again that Fa is classically true, but Fb is not classically true, and $aI_F b$. Then from the reflexivity of indiscriminability $aI_F a$, and so from the base clause for tolerant truth, Fa is tolerantly true. But from the base clause for strict truth, Fa is not strictly true, and so from the recursive clause for negation $\neg Fa$ is tolerantly true. Hence, tolerant truth admits of gluts, where both a vague predication and its negation are true.

Note that everything which is strictly true in an interpretation is also classically true in that interpretation, and everything which is classically true in an interpretation is also tolerantly true in that interpretation.[14] The proof is by simultaneous induction using the mutually recursive definition. For the base case, suppose Fa is strictly true. Then from the reflexivity of indiscriminability and the base clause for strict predication it follows that Fa is also classically true. Likewise, suppose Fa is classically true. Then from the reflexivity of indiscriminability and the base clause for tolerant predication, it follows that Fa is also tolerantly true.

For the inductive step we consider only negation, since conjunction, disjunction and quantification are all straightforward. First suppose that $\neg A$ is strictly true. From the recursive clause for negation, it follows that A is not tolerantly true. From the inductive hypothesis, it follows A is not classically true. So $\neg A$ is classically true. Likewise, suppose that $\neg A$ is classically true. It follows that A is not classically true. From the inductive hypothesis, it follows that A is not strictly true. But then from the recursive clause for negation, it follows that $\neg A$ is tolerantly true.

Corresponding to the three notions of truth, we can define three types of tautology – a sentence is a strict tautology if and only if it is strictly true in every interpretation, a classical tautology if and only if it is classically true in every interpretation, and a tolerant tautology if and only if it is tolerantly true in every interpretation. So, for example, it follows that $Fa \vee \neg Fa$ is a classical and tolerant tautology, but not a strict tautology, since $Fa \vee \neg Fa$ is classically and tolerantly but not strictly true in every interpretation.

Tolerance is not a classical tautology. For consider again an interpretation in which Fa is classically true, but Fb is not classically true, and aI_Fb. Then $(Fa \wedge aI_Fb) \rightarrow Fb$ is not classically true, since $Fa \wedge aI_Fb$ is classically true and Fb is classically false. Likewise, tolerance is not a strict tautology. For consider an interpretation in which Fa is classically true, Fb is classically true, but Fc is classically false, and in which aI_Fb, bI_Fc, but not aI_Fc. Then

[14]This is lemma 1 in Cobreros et al. 2012, p. 357.

$(Fa \wedge aI_Fb) \rightarrow Fb$ is not strictly true, since $Fa \wedge aI_Fb$ is strictly true, but Fb is not strictly true.

However, tolerance *is* a tolerant tautology.[15] Recall that from the definition of the conditional, $(Fa \wedge aI_Fb) \rightarrow Fb$ is equivalent to $\neg(Fa \wedge aI_Fb) \vee Fb$. From the usual clause for disjunction, this is tolerantly true if and only if $\neg(Fa \wedge aI_Fb)$ is tolerantly true or Fb is tolerantly true. Now consider two cases – either Fb is classically true, or Fb is not classically true. In the first case, if Fb is classically true, then from the reflexivity of the indiscriminability relation and the base clause for tolerant truth, Fb is tolerantly true, so $(Fa \wedge aI_Fb) \rightarrow Fb$ is tolerantly true as well.

In the second case, suppose Fb is not classically true. From the recursive clause for negation, $\neg(Fa \wedge aI_Fb)$ is tolerantly true if and only if $Fa \wedge aI_Fb$ is not strictly true. Now either aI_Fb or not aI_Fb. If the former, then since aI_Fb and Fb is not classically true, it follows from the base clause for strict truth that Fa is not strictly true and so $Fa \wedge aI_Fb$ is not strictly true. But if the latter, $Fa \wedge aI_Fb$ is not strictly true because aI_Fb is not true. Either way, $\neg(Fa \wedge aI_Fb)$ is tolerantly true, so in the second case $(Fa \wedge aI_Fb) \rightarrow Fb$ is tolerantly true as well.

Since this argument works for any a and b, it follows that every instance of $(Fx \wedge xI_Fy) \rightarrow Fy$ is tolerantly true, and so $(\forall x)(\forall y)((Fx \wedge xI_Fy) \rightarrow Fy)$, or in other words, tolerance is tolerantly true in every interpretation. So in this framework, there is at least one sense in which tolerance is not only true, but a tautology – viz., although tolerance is not a strict or classical tautology, it is a tolerant tautology. This suggests that proponents of tolerance could helpfully explicate their position in terms of tolerant truth. In the following sections, we explore the implications of this explication of tolerance for the traditional and distributed sorites paradoxes.

[15]See van Rooij 2011, p. 213 and Cobreros et al. 2012, p. 354 for this point.

5. Classical and Three-Valued Logics

Addressing the sorites paradox requires not just an account of truth, but also an account of validity. Continuing to think of validity as preservation of truth, one may define nine kinds of validity, corresponding to the three times three options for the truth of the premises and conclusion (Cobreros et al. 2012, p. 366). In this section, we set the stage for the strict-tolerant approach by briefly considering the three simplest options: defining validity as preservation of classical, strict, or tolerant truth. As it turns out, each of these options lead to strict conservative extensions of three more familiar logics for modelling vagueness.

In the first case, we have: an argument is valid if and only if necessarily, if the premises are classically true, then the conclusion is also classically true. This leads to the logic known as *cc*, for "classical-classical", which turns out to be a strict and conservative extension of classical logic. The extension is strict, since for example the reflexivity and symmetry of indiscriminability imply that aI_Fa is a tautology and $aI_Fb \therefore bI_Fa$ is valid, but conservative since it includes no additional tautologies or valid arguments couched only in the usual vocabulary of classical logic.

Although *cc* is a strict conservative extension of classical logic, implicit sorites arguments are still invalid, for any $n \geq 1$. To see why, just consider an interpretation in which Fa_0 is classically true, Fa_n is not classically true, $a_mI_Fa_{m+1}$ is true for all m such that $0 \leq m < n$, and in which there are no other indiscriminabilities except those required by symmetry and reflexivity. Then the premises of the implicit sorites argument are all classically true, while the conclusion is not classically true, and so the implicit sorites is invalid in *cc*, even for a single step or $n = 1$. But of course, explicit sorites arguments are still valid in *cc* simply because it is an extension of classical logic, and so proponents of *cc* must still deny the premise of tolerance.

For exactly the same reason, implicit distributed sorites are invalid in *cc* for any $n \geq 1$. To see why, consider the same countermodel as in the previous paragraph, but also let $f(a_m) = a_{m-1}$ for all m such that $0 < m \leq n$, and let $f(a_0) = a_n$. Then the premises of the implicit distributed sorites are all classically true, while the conclusion is classically false, and so the implicit distributed sorites is invalid in *cc*, even for a single-step or $n = 1$. But of course, explicit sorites arguments are still valid in *cc* simply because it is an extension of classical logic, and so proponents of *cc* must still deny the premise of tolerance to resolve the explicit distributed sorites too.

In the second case we have: an argument is valid if and only if necessarily, if the premises are strictly true, then the conclusion is strictly true as well. This leads to the logic known as *ss*, for "strict-strict", which turns out to be a strict and conservative extension of the well-known gappy three-valued logic K_3, in which the third-value is interpreted as neither true nor false.[16] As in *cc*, explicit but not implicit, sorites arguments are valid in *ss*, and so the proponent of *ss* must also join the classical logician in rejecting tolerance as a premise of the sorites argument, so we do not consider it further.

In the third case we have: an argument is valid if and only if necessarily if the premises are tolerantly true, then the conclusion is also tolerantly true. This leads to the logic known as *tt*, for "tolerant-tolerant", which turns out to be a strict conservative extension of the well-known glutty three-valued logic *LP*, in which the third value is interpreted as *both* true and false.[17] Since tolerance is tolerantly true in every interpretation, tolerance is a tautology of *tt*, so unlike proponents of classical logic, *cc*, or *ss*, proponents of *tt* cannot resolve the traditional or the distributed sorites paradox by denying tolerance as a premise.

---

[16] K_3 is named for Kleene 1952, p. 334. Tye 1990, p. 544 applies K_3 to vagueness.

[17] Asenjo 1966 and Priest 1979 introduced *LP*. Colyvan 2008, Priest 2010, p. 73 and Ripley 2013b apply *LP* to vagueness. Weber 2010, pp. 1043–5 applies the relevant logic *DK*, which extends *LP* with a relevant conditional. Beall 2014 criticises applying *LP* to vagueness; Weber et al. 2014 reply.

But proponents of *tt* also need not deny tolerance, since even an explicit one-step sorites is invalid in *tt*. For similar reasons, modus ponens too is invalid in *tt*.[18] (This is a manifestation of the invalidity of modus ponens in *LP*, assuming the conditional in *LP* is also defined as usual.[19]) Abandoning modus ponens is one route to resolving the sorites paradox, and could of course solve the distributed sorites too. Nevertheless, whatever the merits of a solution to the sorites paradox which preserves tolerance while abandoning modus ponens, it is of no help to the proponent of noniteration, who wants also to preserve at least one step of soritical reasoning.

6. Strict-Tolerant

Alternatively, we could also define validity thus: an argument is valid if and only if necessarily if the premises are *strictly* true, then the conclusion is *tolerantly* true. This leads to the logic known as *st*, for "strict-tolerant".[20] If we think of a tautology as a sentence which validly follows from no premises, then the tautologies of *st* are just the tolerant tautologies, and so tolerance itself is a tautology of *st*, just as it is of *tt*. But *st* is more suitable than *tt* for modelling noniteration, because it validates both modus ponens as well as one- and two-step sorites arguments, but invalidates implicit sorites arguments of three-steps or more.[21]

[18] See van Rooij 2011, p. 213 and Cobreros et al. 2012, p. 373 for the invalidity of modus ponens in *tt*.

[19] Priest 2010, pp. 73–4, Ripley 2013b, pp. 346–7, and Weber 2010, p. 1040 all allow that *LP* may be extended with conditionals which do validate modus ponens. They concede sorites arguments restated with these conditionals are valid, but deny they are sound.

[20] We also have *sc*, *ct*, *cs*, *tc* and *ts*, for "strict-classical", "classical-tolerant", "classical-strict", "tolerant-classical" and "tolerant-strict" respectively for the five remaining logics (Cobreros et al. 2012, p. 366).

[21] See Cobreros et al. 2012, p. 376. The logics *ct* and *sc* also validate single-step sorites, but invalidate sorites arguments of two-steps or more. Cobreros et al. 2012, p. 373 prefer *st* to *sc* because tolerance is not a tautology of the latter. Although van Rooij 2011, p. 213 initially singled-out *ct* as an appropriate logic for vagueness, Cobreros et al. 2012, p. 373

To get an intuitive feel for validity in st, note that premises are held to a stricter standard than conclusions. In a one step implicit sorites, you can fall from strictly true premises to a classically true conclusion. In a two step implicit sorites, you can fall from strictly true premises to a tolerantly true conclusion – but this is still good enough. However, in a three step implicit sorites, you can fall all the way from strictly true premises to a conclusion that is not even tolerantly true – and this is not good enough. So a pattern of reasoning in st which is reliable enough for one or two steps, may not be reliable enough for three or four.

More precisely, for $n = 2$ or, in other words, the two step-implicit sorites we have $Fa_0, a_0I_Fa_1, a_1I_Fa_2 \vDash_{st} Fa_2$, because if Fa_0 is strictly true and $a_0I_Fa_1$, then from the base clause for strict truth it follows that Fa_1 must be classically true. But if Fa_1 is classically true and $a_1I_Fa_2$, then from the base clause for tolerant truth, Fa_2 must be tolerantly true. Hence, if $Fa_0, a_0I_Fa_1, a_1I_Fa_2$ are all strictly true, then Fa_2 is tolerantly true. (For $n = 1$ we have $Fa_0, a_0I_Fa_1 \vDash_{st} Fa_1$, since if Fa_0 is strictly true and $a_0I_Fa_1$, then Fa_1 is classically true, and then from the base clause for tolerant truth and the reflexivity of indiscriminability, Fa_1 is also tolerantly true.)

However, implicit sorites of $n \geq 3$ or at least three steps are invalid. For a countermodel, let Fa_0 and Fa_1 be classically true, but Fa_{n-1} and Fa_n be classically false, and let it be that $a_mI_Fa_{m+1}$ for all m such that $0 \leq m < n$, but that there are no other indiscriminabilities except those required by reflexivity and symmetry. Then it follows that Fa_0 is strictly true, since only $a_0I_Fa_0$ and $a_0I_Fa_1$, and both Fa_0 and Fa_1 are classically true. Nevertheless, Fa_n is not even tolerantly true, since only $a_nI_Fa_n$ and $a_nI_Fa_{n-1}$, and Fa_n and Fa_{n-1} are classically false, and so the argument is invalid in st.

---

prefer st because it satisfies the deduction theorem. But most of the points in this paper would go through with only superficial changes if ct were adopted instead of st.

Notice that this countermodel shows not only that implicit sorites of at least three steps are invalid, but also that the premises of an implicit sorites of at least three steps are satisfiable which, in *st*, means that they can all be strictly true, and so can be accepted as premises by the proponent of *st*. This point will be more important in section 7, since there I will show that proponents of *st* may resolve the implicit distributed sorites by accepting that it is both valid and sound. Nevertheless, in order to complete this resolution, proponents of *st* must still show that they may accept the premises of the implicit distributed sorites as strictly true or, in other words, that the premises are satisfiable in *st*.

Two features of the strict-tolerant solution to the sorites paradox need to be emphasised. First, *explicit* sorites arguments, in which tolerance is stated as a premise, are still valid in *st* for any number of steps.[22] To see this, recall from section 4 that everything which is strictly true in an interpretation is also classically true in that interpretation, and everything which is classically true in an interpretation is also tolerantly true in that interpretation. It follows that every argument valid in *cc* is also valid in *st*, since if the premises of the argument are strictly true, then they are classically true, and so the conclusion is classically true, and so also tolerantly true.

In other words, *st* is an extension of *cc*. (Moreover, the extension is strict because, for example, implicit sorites of $n \leq 3$ are valid in *st* but not in *cc*.) But as we noted in section 5, explicit sorites arguments are valid in *cc* for any n, and so it follows that explicit sorites arguments are also valid in *st* for any n. Thus, the proponent of *st* must join proponents of *cc* and classical logic in resolving the paradox posed by explicit sorites arguments by rejecting tolerance as a premise. Since the strict-tolerant approach was motivated in large part by the desire to preserve tolerance as an intuitive principle, on the face of it this looks like a major set-back.

---

[22]As Cobreros et al. 2012, p. 376 admit.

Moreover, because tolerance is a tautology of *st*, it would seem as if explicit sorites of any number of steps are not only valid, but also *sound.* If so, then the sorites paradox is back with a vengeance, since without the option of rejecting tolerance as a premise, or denying one of the other premises, we would appear to have no option but to embrace the absurd conclusion. If this were right, then the strict-tolerant approach would not only fail to improve on the classical solution to the problem, but would make the problem worse by closing off that route to a solution altogether.

But this objection is too fast, since it overlooks that the definition of soundness in *st* should be tailored to reflect the definition of validity, thus: an argument is sound in *st* if and only if it is valid in *st* and the premises are *strictly* true. As Cobreros et al. 2012, p. 377 write:

> This version of the paradox is *st*-valid. After all, it is classically valid, and, as we have seen, *st* is stronger than classical logic. Here the reason we do not conclude that Pa_n holds, even tolerantly, is because there is an untrue premise. The third premise, tolerance, is not strictly true, and it is strict truth we require of our premises in *st*.

So although tolerance is a tautology of *st*, it cannot be used as a premise in the sorites argument, because it is merely tolerantly and not strictly true.

The fact remains that proponents of *st* join proponents of classical logic in denying tolerance as a premise of the argument. Nevertheless, proponents of *st* can still style themselves as defenders of tolerance in a way that proponents of classical logic cannot, since while they must reject tolerance as a premise of sorites arguments, because it is not strictly true, they can – and moreover must – continue to accept it as a tautology and a conclusion, because it is tolerantly true. So there is a sense in which proponents of *st* must reject tolerance. But their position is still an improvement, because there is another sense in which they accept tolerance after all.

The second feature that deserves emphasis is that the fact that both a sentence and its negation can be tolerantly true is essential to this solution to the sorites. One way this manifests itself is that even though implicit sorites argument of $n \geq 3$ are invalid in *st*, the following argument *is* valid:

(10) $Fa_0 \wedge \neg Fa_n$

(11) $a_0 I_F a_1, a_1 I_F a_2, ..., a_{n-1} I_F a_n$

(12) $\therefore (\exists x)(Fx \wedge \neg Fx)$

Here, the premises are the same as the premises of an implicit sorites argument, except the first premise is strengthened with the negation of the conclusion. The new conclusion is an explicit inconsistency.

To get an intuitive feel for why this is so, consider that in proceeding via a sorites argument from Abelard, who is strictly tall, to Zach, who is strictly not tall, one must at some point pass through a borderline case of someone who is not strictly tall nor strictly not tall – in other words, a case of someone who is tolerantly tall and tolerantly not tall. Since although premises in *st* must be strictly true, conclusions need only be tolerantly true, this borderline case will be a witness to the tolerant truth of the conclusion that someone is tall and not tall. In general, strict soritical premises in *st* will license inconsistent conclusions about borderline cases.

More precisely, consider an interpretation in which premises 10 and 11 are strictly true. Then from the recursive clause for conjunction and premise 10, it follows that Fa_0 is strictly true and $\neg Fa_n$ is strictly true. Since everything strictly true is also classically true, it follows that Fa_0 is classically true and $\neg Fa_n$ is also classically true. So there must be some m such that $0 \leq m < n$ and Fa_m is classically and so tolerantly true, but Fa_{m+1} is not classically true. But from premise 11, $a_m I_F a_{m+1}$, so Fa_m is not strictly true and $\neg Fa_m$ is tolerantly true, so $Fa_m \wedge \neg Fa_m$ is tolerantly true, as is $(\exists x)(Fx \wedge \neg Fx)$.

Although this version of the sorites is valid for any n, note that the premises are satisfiable in *st* for and only for $n \geq 3$, by the same reasoning we just used to show the traditional implicit sorites is invalid for and only for $n \geq 3$. So as we will see in section 7, there is an exact analogy between this version of the sorites and the implicit distributed sorites, which is valid in *st* for any n, but for which its premises are satisfiable in *st* for and only for $n \geq 3$. Thus, as I will argue in section 7, the resolution of the traditional sorites paradox in *st* extends to a solution of the distributed sorites paradox too.

Note also that this conclusion commits the strict-tolerant approach not to paraconsistency, according to which not everything follows from a contradiction, but to dialetheism proper, according to which some contradictions are true (in this case, as conclusions of arguments).[23] In fact, *st* is not a paraconsistent logic at all since, because a contradiction cannot be strictly true, it validates the inference rule of explosion, viz.: $A, \neg A \vdash B$, and so according to it, everything does follow from a contradiction. Nevertheless, *st* escapes triviality because although it allows tolerantly true contradictions to occur in the conclusions of arguments, it forbids them from recurring as strictly true premises.

This explicit commitment to dialetheism is not obviously part of the initial motivation for noniteration, and so could be seen as a cost of modelling noniteration with *st*. On the other hand, if one sees independent motivation for an approach to vagueness according to which vague predicates and their negations are both true of borderline cases, one may see the ability of this approach to modelling noniteration as a further advantage. In either case, we will see in the following section that the commitment to dialetheism is exactly what enables the strict-tolerant solution to the traditional sorites paradox to resolve the distributed sorites paradox too, at no further cost.

7. Embracing Inconsistency

How should a proponent of the strict-tolerant approach to the sorites paradox respond to the distributed sorites? Whereas implicit traditional sorites

[23] Restall 1997, for example, helpfully distinguishes paraconsistency from dialetheism.

of $n \geq 3$ are invalid in st, implicit *distributed* sorites are valid for any n. To see this, suppose the premises of the implicit distributed sorites are strictly true. From premise 8, Fa_0 and $\neg Fa_n$ are strictly and so also classically true, and $f(a_0) = a_n$. Since n is finite, $f(a_0) = a_n$ and f is a bijection, we can keep applying f to construct a series $a_n, f(a_n), f(f(a_n)), ..., f^k(a_n) = a_0$ for some $k \leq n$, which begins with a_n and eventually terminates in a_0.

Since Fa_0 is classically true but Fa_n is not, there must be some $0 < l \leq k$ in this series such that $Ff^{l-1}(a_n)$ isn't classically true, but $Ff^l(a_n)$ is classically true. Then from the reflexivity of indiscriminability, it follows that $Ff^l(a)$ is tolerantly true. And since $f^{l-1}(a_n) \neq a_0$ it follows from premise 6 that $f^{l-1}(a_n)I_F f^l(a_n)$, so $Ff^l(a_n)$ is not strictly true and so, from the recursive clause for negation, $\neg Ff^l(a_n)$ is tolerantly true. So $Ff^l(a_n) \wedge \neg Ff^l(a_n)$ is tolerantly true, and it follows that $(\exists x)(Fx \wedge \neg Fx)$ is tolerantly true as well, establishing the validity of implicit distributed sorites arguments in st.

So the proponent of st is committed to accepting the validity of implicit distributed sorites arguments of any length, and so committed to accepting the inconsistent conclusion $(\exists x)(Fx \wedge \neg Fx)$ when the premises of the implicit distributed sorites are tolerantly true. However, accepting this inconsistency comes at no further cost, since as we just saw in the previous section, the argument from the premises of the traditional implicit sorites to the conclusion $(\exists x)(Fx \wedge \neg Fx)$ is already valid in st. So the implicit distributed sorites poses no additional problem to the proponent of st.

To put this point in terms of the original example, we can conclude that there is a step on the staircase which both does and does not contain a heap. So when we draw conclusions about the number of heaps on the stairs, we should both count this step as a heap and not count this step as a heap. So we can draw inconsistent conclusions about the number of heaps on the stairs – for some n, the number will be both n and $n + 1$. But whereas in classical logic this would be a reductio of our premises, here it is not, since inconsistent conclusions can be tolerantly true. The solution is not to reject the reasoning, but to accept the inconsistency.

In order to complete their resolution of the implicit distributed sorites paradox, proponents of *st* need to show not only that they can accept that the inconsistent conclusion is tolerantly true, but also that they can accept that the premises are strictly true. In other words, they need to show that the premises of the implicit distributed sorites paradox are *satisfiable*, where a set of sentences is satisfiable in *st* if and only if there is some interpretation according to which they are all *strictly* true.[24] (To see why, think of the set of sentences as being the premises of an argument, and unsatisfiability as meaning that no argument with those premises could be sound, or thus that those premises could not be strictly true.)

Just as a traditional two-step sorites argument is valid in *st*, the premises of an implicit distributed sorites for $n = 2$ is *not* satisfiable in *st*. To see why, suppose the premises of the implicit distributed sorites for $n = 2$ are all strictly true. From premise 8, we have that Fa_0 is strictly and so also classically true, that $\neg Fa_2$ is strictly true and so from the recursive clause for negation that Fa_2 is not tolerantly true, and that $f(a_0) = a_2$. From premise 5 and 7, we have that $f(a_2)$ is either a_0 or else a_1 (since if $f(a_2)$ were a_2, then it would follow from premise 5 and $f(a_0) = a_2$ that $a_0 = a_2$).

In the first case, if $f(a_2)$ is a_0, then from premise 6, we have $a_2 I_F a_0$, in which case since Fa_2 is not tolerantly true, Fa_0 is not classically true, contradicting what we said earlier. But in the second case if $f(a_2)$ is a_1, then from premise 6, we have $a_2 I_F a_1$, in which case since Fa_2 is not tolerantly true, Fa_1 is not classically true. Then from premise 5, we have that $f(a_1)$ is a_0. Then from premise 6, we have $a_1 I_F a_0$, in which case since Fa_2 is not classically true, Fa_0 is not strictly true, again contradicting what we said.

However, just as a traditional three-step sorites argument is invalid in *st*, the premises of an implicit distributed sorites for $n \geq 3$ *is* satisfiable in

[24]See Cobreros et al. 2012, p. 357 for their definition of satisfiability.

st. The model is the same as in the previous section: let Fa_0 and Fa_1 be classically true, but Fa_{n-1} and Fa_n be classically false, and let it be that $a_m I_F a_{m+1}$ for all m such that $0 \leq m < n$, but that there are no other indiscriminabilities except those required by reflexivity and symmetry. We only have to add that $f(a_0) = a_n$ and $f(a_m) = a_{m-1}$ for all $0 < m \leq n$.

How about *explicit* distributed sorites, which explicitly include the principle of tolerance? Since explicit distributed sorites are valid in classical logic, and *st* is stronger than classical logic, explicit distributed sorites of any length are valid in *st* too. This of course is no problem for the proponent of *st*, since they have already accepted the conclusion of the explicit distributed sorites argument both in their response to the traditional implicit sorites, and again in their response to the implicit distributed sorites.

However, explicit distributed sorites of any size are not only valid, but also have premises which are unsatisfiable. For recall that everything strictly true in an interpretation is also classically true in that interpretation. But since the premises of the explicit distributed sorites are classically inconsistent, they cannot be classically true in any interpretation, and so cannot be strictly true in any interpretation either. Since the premises of the explicit distributed sorites cannot all be strictly true, the proponent of *st* must reject at least one of them, and so in cases where the other premises are all true must join proponents of classical logic in rejecting tolerance as a premise.

Given that *st* was motivated in part by the desire to preserve tolerance as an intuitive principle, this may seem strange. But recall from the last section that proponents of the strict-tolerant approach must also reject tolerance as a premise of traditional explicit sorites arguments, which remain valid in *st*. But just as in the previous section, proponents of *st* can still style themselves as defenders of tolerance, since while they must reject it as a premise, they still accept it as a tautology. Insofar as one was satisfied by this defence of tolerance in response to the traditional sorites, one should be equally satisfied with the same defence in response to the distributed sorites.

So like the strict-tolerant solution to the traditional sorites, the strict-tolerant solution to the distributed sorites involves both an explicit commitment to dialetheism, and a sense in which tolerance is explicitly rejected. Perhaps these features are objectionable. But since they were both already features of the strict-tolerant solution to the traditional sorites paradox, they are not more objectionable now than they were before. Insofar as the strict-tolerant approach succeeded as a solution to the traditional sorites paradox, it succeeds as a solution to the distributed sorites paradox too.[25]

References

Asenjo, F. G. (1966). "A Calculus of Antinomies." *Notre Dame Journal of Formal Logic* 7.1, pp. 103–105. DOI: 10.1305/ndjfl/1093958482.

Barnett, Zach (2018). "No Free Lunch: The Significance of Tiny Contributions". *Analysis* 78.1, pp. 3–13. DOI: 10.1093/analys/anx112.

– (2019). "Tolerance and the Distributed Sorites". *Synthese* 196.3, pp. 1071–1077. DOI: 10.1007/s11229-017-1499-4.

Beall, J. C. (2014). "Finding Tolerance without Gluts". *Mind* 123.491, pp. 791–811. DOI: 10.1093/mind/fzu081.

Beall, J. C. and Greg Restall (2006). *Logical Pluralism.* Oxford: Oxford University Press.

[25]The idea for this paper arose from a discussion with my colleagues Zach Barnett and Lavinia Picollo; I'm especially grateful to Zach for much discussion. I'm also grateful to the students in my non-classical logic class, Eugene Ho, Jethro Peck, Ashish Stephen Peter, and Thomas Tan for working through the material with me. I thank my colleagues Zach Barnett, Ethan Jerzak, Michael Pelczar, Hsueh Qu, Mattias Skipper, Lavinia Picollo, Weng-Hong Tang and Daniel Waxman for discussing the paper in our reading group. Finally, I thank Nathaniel Gan for reading two recent drafts of the paper, and anonymous referees for their comments. This paper was part of our Metaphysics of Humanity project, supported by the Ministry of Education, Singapore, grant number A-0003057-00-00.

Carlson, Erik, Magnus Jedenheim-Edling, and Jens Johansson (2021). "The Significance of Tiny Contributions: Barnett and Beyond". *Utilitas* 33.4, pp. 488–496.

Cobreros, Pablo (2011). "Paraconsistent Vagueness: A Positive Argument". *Synthese* 183.2, pp. 211–227. JSTOR: 41477660.

Cobreros, Pablo, Paul Egré, Robert van Rooij, and David Ripley (2012). "Tolerant, Classical, Strict". *Journal of Philosophical Logic* 41.2, pp. 347–385. DOI: 10.1007/s10992-010-9165-z.

Colyvan, Mark (2008). "Vagueness and Truth". *From Truth to Reality*. Ed. by Heather Dyke. New York: Routledge, pp. 29–40.

Égré, Paul (2015). "Vagueness: Why Do We Believe in Tolerance?" *Journal of Philosophical Logic* 44.6, pp. 663–679. DOI: 10.1007/s10992-015-9352-z.

Gaifman, Haim (2010). "Vagueness, Tolerance and Contextual Logic". *Synthese* 174.1, pp. 5–46. DOI: 10.1007/s11229-009-9683-9.

Hyde, Dominic (1997). "From Heaps and Gaps to Heaps of Gluts". *Mind* 106.424, pp. 641–660. JSTOR: 2254669.

– (2010). "The Prospects of a Paraconsistent Response to Vagueness". *Cuts and Clouds*. Ed. by Richard Dietz and Sebastiano Moruzzi. Oxford University Press. DOI: 10.1093/acprof:oso/9780199570386.001.0001.

Hyde, Dominic and Mark Colyvan (2008). "Paraconsistent Vagueness: Why Not?" *The Australasian Journal of Logic* 6. DOI: 10.26686/ajl.v6i0.1798.

Hyde, Dominic and Diana Raffman (2018). "Sorites Paradox". *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Summer 2018. Metaphysics Research Lab, Stanford University.

Kleene, Stephen (1952). *Introduction to Metamathematics*. Groningen: Wolters-Noordhoff.

Priest, Graham (1979). "The Logic of Paradox". *Journal of Philosophical Logic* 8.1, pp. 219–241. DOI: 10.1007/BF00258428.

– (2006). *In Contradiction: A Study of the Transconsistent*. 2nd Edition. Oxford: Oxford University Press.

– (2010). "Inclosures, Vagueness, and Self-Reference". *Notre Dame Journal of Formal Logic* 51.1, pp. 69–84. DOI: 10.1215/00294527-2010-005.

Restall, Greg (1997). "Paraconsistent Logics!" *Bulletin of the Section of Logic of the Polish Academy of Sciences* 26.3, pp. 156–163.

Ripley, David (2013a). "Paradoxes and Failures of Cut". *Australasian Journal of Philosophy* 91.1, pp. 139–164. DOI: 10.1080/00048402.2011.630010.

– (2013b). "Sorting out the Sorites". *Paraconsistency: Logic and Applications*. Ed. by Koji Tanaka, Francesco Berto, Edwin Mares, and Francesco Paoli. Dordrecht: Springer Netherlands, pp. 329–348. DOI: 10.1007/978-94-007-4438-7_18.

Sorensen, Roy (1988). *Blindspots*. Oxford: Oxford University Press.

Tye, Michael (1990). "Vague Objects". *Mind* 99.396, pp. 535–557. JSTOR: 2255019.

Van Rooij, Robert (2011). "Vagueness, Tolerance and Non-Transitive Entailment". *Understanding Vagueness. Logical, Philosophical and Linguistic Perspectives*. Ed. by Petr Cintula, Christian G. Fermüller, and Petr Hájek. College Publications, pp. 205–221.

Weber, Zach (2010). "A Paraconsistent Model of Vagueness". *Mind* 119.476, pp. 1025–1045. DOI: 10.1093/mind/fzq071.

Weber, Zach, David Ripley, Graham Priest, Dominic Hyde, and Mark Colyvan (2014). "Tolerating Gluts". *Mind* 123.491, pp. 813–828. DOI: 10.1093/mind/fzu057.

Zardini, Elia (2008). "A Model of Tolerance". *Studia Logica* 90.3, pp. 337–368. DOI: 10.1007/s11225-008-9156-z.