

Chapter 10

Tolerance and Degrees of Truth



Pablo Cobrerros, Paul Égré, David Ripley, and Robert van Rooij

Abstract This paper explores the relations between two logical approaches to vagueness: on the one hand the fuzzy approach defended by Smith (*Vagueness and degrees of truth*. Oxford University Press, Oxford, 2008), and on the other the strict-tolerant approach defended by Cobrerros et al. (“Tolerant, Classical, Strict”, *J Philos Logic* 41(2):347–385, 2012). Although the former approach uses continuum many values and the latter implicitly four, we show that both approaches can be subsumed under a common three-valued framework. In particular, we defend the claim that Smith’s continuum many values are not needed to solve what Smith calls ‘the jolt problem’, and we show that they are not needed for his account of logical consequence either. Not only are three values enough to satisfy Smith’s central desiderata, but they also allow us to internalize Smith’s closeness principle in the form of a tolerance principle at the object-language. The reduction, we argue, matters for the justification of many-valuedness in an adequate theory of vague language.

P. Cobrerros
University of Navarra, Pamplona, Spain
e-mail: pcobrerros@unav.es

P. Égré (✉)
IRL Crossing, CNRS, Adelaide, Australia – ENS-PSL University, Paris, France
e-mail: paul.egre@ens.fr

D. Ripley
Monash University, Clayton, VIC, Australia
e-mail: david.ripley@monash.edu

R. van Rooij
ILLC, Amsterdam, The Netherlands
e-mail: R.A.M.vanRooij@uva.nl

10.1 Introduction

In this paper, we explore the relations between two logical approaches to vagueness: the degree-theoretic approach of Smith (2008) and the strict-tolerant account of vagueness originally laid out in Cobreros et al. (2012). At first glance, these approaches look quite different: the former is based on a fuzzy logic, that is on continuum many truth values, whereas the latter is implicitly based on a four-valued logic. Conceptually, the approaches differ still further: on Smith's approach, the principle of *tolerance* is rejected (the principle whereby anyone imperceptibly shorter than a tall person must count as tall), to be replaced with a weaker principle of *closeness* (according to which for any two persons x and y with imperceptibly distinct heights, the semantic values of “ x is tall” and “ y is tall” must be close). Smith moreover articulates both principles as metalinguistic constraints. On the strict-tolerant logic of vagueness developed in Cobreros et al. (2012), on the other hand, the principle of tolerance is stated in the object-language, and it is valid.

Our point of departure in this paper is that there is less to these differences than meets the eye. The two approaches have much in common. In particular, they can both be cast into a common three-valued framework (Cobreros et al., 2015b), as we proceed to explain below. This is because the accounts both work with a common structure: they have it that whenever all the premises of an argument hold to some particular strong standard, the conclusion must hold to some particular weaker standard (a form of what we call “permissive consequence” in Cobreros et al., 2015b). It is this common structure that is directly captured by a three-valued framework. Moreover, despite the kinds of models these approaches use, both stick to full first-order classical logic, validating every classically-valid argument. Cobreros et al. (2012) but not Smith (2008) extends the object language to allow for the principle of tolerance to be stated.¹

In what follows, we consider extended versions of both approaches, in order to bridge the gap between them. One of our goals in this paper is logical: we introduce a family of consequence relations, which we call *parameterized consequence* relations, which basically subsume both Smith's consequence and strict-to-tolerant consequence (abbreviated ST-consequence) as particular cases. We use the framework to show that strict-to-tolerant consequence in a sense occupies a unique position among this family of consequence relations when the language is sufficiently expressive to accommodate tolerance principles for vague predicates. Our main goal, however, is more philosophical, and concerns the role and the number of truth values in an adequate theory of vague language.

Two roles are sometimes distinguished for truth values: a *referential* role, used to assign values to sentences, and a logical or *inferential* role, concerned with entailment relations between sentences (Suszko, 1977; Malinowski, 2009b; Shramko and Wansing, 2011). Smith's theory and the strict-tolerant theory do not

¹ Strict-tolerant logic conservatively extends classical logic (it is fully classical over its \sim -free fragment), but it is nontransitive over its \sim -full fragment. The loss of transitivity is arguably a non-classical feature of ST-logic. See Cobreros et al. (2012) and below for discussion.

use the same number of truth values to serve as references for sentences, but our argument will show that the richness of Smith's referential apparatus can be cut down to three values when it comes to capturing central inferential principles about vagueness. There is a clear and precise sense, therefore, in which the strict-tolerant approach is canonical for the kind of treatment of vagueness advocated by Smith. This reduction is not meant to offer a full-fledged theory of the relation between referential and inferential many-valuedness for vague language, but we take it to be a step in that direction, not least because vagueness is an area in which the introduction of intermediate truth values between 0 and 1 has been a source of controversy and remains in need of further justification (see Williamson, 1994; Haack, 1996; and Smith, 2008 for rival views).

Our paper is structured as follows. The next section gives a brief review and a philosophical discussion of Smith's account of closeness and tolerance, rebutting his arguments against three-valued approaches to vagueness. Section 10.3 gives a more detailed formal comparison between Smith's account and the strict-tolerant account. Section 10.4 introduces the notion of parameterised consequence and shows how to embed both Smith's consequence and ST-consequence under that scheme. In Sect. 10.5, finally, we draw more general lessons from this comparison, in particular regarding the truth values needed to study vagueness.

10.2 Tolerance and Closeness

Vague predicates seem to exhibit a phenomenon known as *tolerance* (Wright, 1975; Kamp, 1981). For example, consider the vague predicate "young": there seem to be differences in age too small to matter for youngness (such as a difference of a nanosecond for age in humans, at least in most contexts). Tolerance is the claim that this seeming is correct: that there really are differences too small to matter. In general, for a predicate P , P -similarity (which we write \sim_P) is a relation that holds between things when, according to tolerance, they are too similar in whatever respects matter for P for it to be the case that P applies to one of them but not the other. The principle of tolerance may then be stated as follows:

- (1) $\forall x \forall y (Px \wedge x \sim_P y \rightarrow Py)$ (Tolerance)

This is a notoriously problematic principle in two-valued classical logic, for conjoined with the existence of a soritical series, it gives rise to the sorites paradox. For vague P , it is easy to imagine a list of individuals, each P -similar to the next, but where the first is obviously P and the last obviously not P . The way this problem is evaded in classical logic usually involves a rejection of the tolerance principle.² Rejecting tolerance in this two-valued setting, however, leaves us with the idea that

² Bare rejection of the tolerance principle, of course, doesn't give much of a useful theory. Because of this, rejection of tolerance is often paired with some explanation of its intuitive appeal. These explanations might involve, for example, epistemology (Sorensen, 2001; Williamson, 1994), context-sensitive concepts (Fara, 2000; Raffman, 1996), or pragmatic restrictions (Manor, 2006; Gaifman, 2010; van Rooij, 2011; Pagin, 2011; Gómez-Torrente, 2011). Tolerance can be

in every sorites series, there are at least two individuals a and b that are very close in P -relevant respects, but where Pa and Pb are still assigned opposite truth values.³

This violation of tolerance leaves us with a sudden jolt in truth values as we proceed along a sorites series; Smith thus calls it ‘the jolt problem’. (On Smith’s view, by contrast, the core of vagueness lies in the absence of such jolts: gradual modifications of the features relevant for a property should be matched by gradual modifications of the truth values assigned to the claim that the property applies.) Consequently, neither tolerance nor its negation can provide a satisfactory account of vagueness in this two-valued setting: tolerance because it leads to paradox, and the negation of tolerance because of the jolt problem.

According to Smith, the proper way to evade this dilemma is to abandon two-valued logic and replace the principle of tolerance by a principle of *closeness*. Given a vague predicate P , closeness states that if two objects a and b are P -similar, then the sentences Pa and Pb should have truth values that are very close to each other. Formally, this may be represented as follows, letting \sim_T stand for closeness between truth values, and $\ll Pa \gg$ for the truth value of Pa :

(2) If $a \sim_P b$, then $\ll Pa \gg \sim_T \ll Pb \gg$ (Closeness)

Closeness is the leading principle behind the introduction of degrees of truth in Smith’s approach. Two degrees of truth seem obviously inadequate to accommodate closeness, so more than two truth values are required. But how many more, if closeness is to be secured? Smith’s answer is as follows:

I do not know exactly how many degrees of truth we need in order to accommodate Closeness. The point is simply that we need a significant number of them. [...] As far as accommodating vagueness goes, we might have a large finite number of degrees of truth [...] or we might have continuum-many degrees of truth (as in the fuzzy picture) (p. 190).

Smith’s choice is to have continuum many degrees of truth, though Smith admits that there are no conclusive reasons against having a finite number of degrees of truth.⁴ An option Smith explicitly rejects, however, is working with only three degrees of truth (Smith, 2008, p. 186; 2005, p. 178). According to Smith, a third-value view will necessarily suffer from the jolt problem. Smith’s reason to reject trivalent approaches is as follows (p. 186):

If one sentence is True and another False, then they are as far apart as can be in respect of truth — and furthermore, they are in an absolute sense *very far apart* in respect of truth. Given that Truth and Falsity are poles apart in this way, no third truth status can be *very close* to *both* of them.

Smith’s argument is meant to be fully general, that is, it ought not to depend on the interpretation of the third truth value. What about the case, however, in which

maintained, at least in some cases, by denying that there can be such a chain of P -similar individuals in the first place; see Fara (2001).

³ The needed quantifier move here—from $\neg\forall$ to $\exists\neg$ —is not valid in intuitionistic logic. We don’t discuss intuitionistic approaches to the sorites here; see Read and Wright (1985), Wright (2001), Rumfitt (2015) for more.

⁴ In particular, Smith rejects arguments based on the seeming arbitrariness of any particular finite number; see p. 190.

the third truth value is interpreted as “Both-True-and-False”? (For approaches to vagueness that take this route, see Cobreros (2013), Cobreros (2013), Priest (2013), Priest (2013), Ripley (2013b), Ripley (2013b), Weber et al. (2014), Weber et al. (2014).) Isn’t this third truth value close to “True” and close to “False”?

Take two P -similar objects a and b such that Pa gets value 1, and Pb the value 0.5: since this means that Pa is True, and that Pb is True and False, there no longer is any jolt in this case, since the value assigned to Pb retains some element of the value assigned to Pa . These claims match in an important respect; they are both true. The same holds if Pc is assigned 0.5 and Pd is assigned 0. These claims match in an important respect; they are both false. The only jolt would be a situation in which a and b are P -similar, but such that Pa gets 1 and Pb 0—and this is exactly what does not arise, on such an approach. So, arguably, a glut-based theory of vagueness does not predict any jolts.⁵

This argument would not suffice to convince Smith, however. For Smith’s point is that regardless of its dialetheist definition, a value such as Both-True-and-False cannot be *very close* both to True and to False. Thus, Smith writes (p. 186):

For if one thing is *very* similar to each of two other things in some respect, then those two things must at the very least be *reasonably* similar to one another in that respect — yet Truth and Falsity are not similar at all in respect of truth. Thus, to the extent that a sentence is very close to True, it is not very close to False, and vice versa.

Our response here is that True and False can actually be taken to be “reasonably similar”, though we need to be careful about the relevant respect. While Smith writes that they are “not similar at all in respect of *truth*”, we can observe that they are similar in respect of being *truth-values*, consistently with being as far apart as can be along that dimension. Consider the following analogy: “black” and “white” are certainly not similar at all “in respect of (their proximity to) black”, but they are similar in respect of being achromatic colors. And in that respect black and white are indeed reasonably similar. And while they are poles apart along that dimension of comparison, grey is a mixture which is as close to white and to black as can be. Moreover, because no property other than grey can be close to both white and black while being distinct from either, grey is even *very close* to both.

Admittedly, a central grey, say of RGB value (128, 128, 128), may be seen as relatively distant from a central white (255, 255, 255) and from a central black (0, 0, 0). But in our analogy we need not equate “grey” with a specific triple of RGB values. Consider the region of triples of form (x, x, x) with $0 < x < 255$. This is the grey region (including dark and light greys), and it is very similar to the (just) white region, and very similar to the (just) black region. Structurally, our point is that “Both-True-and-False” may be viewed in the same way, as denoting a region of overlap.⁶

⁵ See Égré (2011b) for a preliminary version of this objection. Here, however, we focus on Smith’s topological argument against three values. We thank an anonymous for urging us to do so.

⁶ This analogy can be made rigorous: Égré (2021) treats “true” and “false” as *absolute* gradable predicates, structurally ambiguous between a *total* and a *partial* interpretation (in the sense of

Statistical theories of vagueness naturally fit with this interpretation of the third value, but this interpretation too does not give rise to jolts (see Borel, 1907; Égré, 2011a; Lassiter, 2011; Égré & Barberousse, 2014; Lassiter & Goodman, 2017; Égré, 2017). On such an approach, we can view 1 as applying to an item for which the response can only be of the form “ x is P ”, 0.5 as encoding an item for which responses can be either “ x is P ” or “ x is not P ”, and 0 as encoding an item for which the response can be only of the form “ x is not P ”. (Of course, items that get assigned the value 0.5 on this picture may still be such as to support different proportions or probabilities of “ P ”-responses over “not P ”-responses. But that is not to say that one should necessarily semanticize those different proportions into distinct truth values.) On this interpretation, the last item that gets the value 1 in a sorites and the first that gets the value 0.5 still have no jolt between them, in that they warrant identical responses on most occasions.

What the glutty and statistical views have in common is that the value 0.5 is seen as encoding a way in which the statuses represented by 1 and 0 *interact*, rather than encoding a distinct status. This makes them altogether different from an interpretation on which 0.5 encodes a distinct response from the ones encoded by 1 or 0, such as “I don’t know” or “Indeterminate”. These third-status interpretations predict two P -similar (and so, consecutive in a sorites series) items that *mandate* distinct responses, instead of *allowing* for identical responses. It is only these third-status interpretations that fall victim to Smith’s jolt problem (For related discussion, see Wright, 2001, 2003). Smith’s argument, then, isn’t fully about the *number* of values in play, but instead turns also on their *interpretation*.

Consequently, while we agree with Smith that a theory of vagueness working with only two values will incur the jolt problem, our point is that three truth values can suffice to give adequate provision against it. Like Smith, we also agree that the following version of Tolerance (stated p. 160), which we call Smith-tolerance, should be rejected:

(3) If $a \sim_P b$ then $\llbracket Pa \rrbracket = \llbracket Pb \rrbracket$ (Smith-Tolerance)

But unlike Smith, we think that the tolerance principle in the version stated in (1) can be preserved as a first-order claim in the object language of a theory of vagueness (see Cobreros et al., 2012). In what follows, we will show that a three-valued version of Closeness is actually enough to support Tolerance as an object-language principle; since it is this object-language formulation that generates the best form of the jolt problem, we conclude that three-valued approaches need not be subject to jolts.

Rotstein and Winter (2004)). The total interpretations denote disconnected endpoints on the scale, but the partial interpretations overlap. In the case of achromatic colors, the total interpretation of “white” in RGB triples is conventionally $\{(255, 255, 255)\}$, but the partial one is the upset $\{(x, x, x) | k \leq x \leq 255\}$, and dually for “black”, its total interpretation is conventionally $\{(0, 0, 0)\}$, and the partial one the downset $\{(x, x, x) | 0 \leq x \leq k'\}$: provided $k < k'$, the partial interpretations overlap.

10.3 Smith-Consequence and ST: A Comparison

In this section, we present and briefly compare two logical systems: the one adopted in Smith (2008), which we will call ‘Smith-consequence’, and a slightly modified version of the one adopted in Cobreros et al. (2012), which we will call ‘ST’. (The purpose of the modifications is twofold: for simplicity, and to ease comparison with Smith’s approach.) At first, we will consider versions of these logics that completely ignore the connection between the object-language predicates P and $\sim P$. As far as we are concerned for now, these are simply distinct predicates, and do not constrain each other in any way. We will introduce connections between these predicates later, once we have described the common structure that underlies both Smith-consequence and ST.

10.3.1 Smith

Smith’s approach to vagueness is based on models that assign values from the real interval $[0, 1]$ to formulas; these models assign values to compound formulas compositionally along the usual Łukasiewicz lines (viz. Łukasiewicz, 1920) for the 3-valued case), *except* for the conditional. Smith leaves out Łukasiewicz’s conditional entirely, instead defining a material conditional $A \rightarrow B$ to be equivalent to $\neg A \vee B$. In sum, then, Smith’s models are pairs $\langle D, I \rangle$ of a domain and an interpretation, where:

- For a term t , $I(t) \in D$
- For an n -ary predicate P , $I(P) \in [0, 1]^{(D^n)}$
- For an atomic sentence $A = P(t_1, \dots, t_n)$, $I(A) = I(P)(I(t_1), \dots, I(t_n))$
- $I(\neg A) = 1 - I(A)$
- $I(A \wedge B) = \min(I(A), I(B))$
- $I(A \vee B) = \max(I(A), I(B))$
- $I(A \rightarrow B) = \max(1 - I(A), I(B))$
- $I(\forall x A(x)) = \text{glb}\{I'(A(x)) : I' \text{ is an } x\text{-variant of } I\}$
- $I(\exists x A(x)) = \text{lub}\{I'(A(x)) : I' \text{ is an } x\text{-variant of } I\}$

These are the models we will work with for the remainder of the paper; we will simply call them *models*. (When it is convenient, for a model $M = \langle D, I \rangle$ we will sometimes write $M(A)$ instead of $I(A)$ for the value of a formula A .)

An *argument*, for our purposes here, is something of the form $\Gamma \succ \Delta$, where Γ and Δ are sets of formulas; this should be thought of as the argument with premises Γ and conclusions Δ . A *consequence relation* is a set \models of arguments; we will write $\Gamma \models \Delta$ for the claim that $\Gamma \succ \Delta \in \models$. To specify a consequence relation \models model-theoretically, we simply specify a relation—the *countermodel* relation—between models and arguments; we then say that $\Gamma \models \Delta$ iff there is no model M such that M is a countermodel to $\Gamma \succ \Delta$.

Although the models we work with are relatively familiar, Smith's approach to consequence is not a usual one for models like these.

Definition 10.1 A model $\langle D, I \rangle$ is a *Smith countermodel* to an argument $\Gamma \succ \Delta$ iff:

- $I(\gamma) > 0.5$ for every $\gamma \in \Gamma$ and
- $I(\delta) < 0.5$ for every $\delta \in \Delta$.

If there is no Smith countermodel to an argument $\Gamma \succ \Delta$, then the argument is *Smith valid*, written $\Gamma \models_S \Delta$.⁷ □

This approach effectively treats the models as having *three* inferential statuses to assign to formulas. A model can assign a formula a value strictly greater than 0.5, suitable to be the value of a premise of an argument in a countermodel to that argument; or a value strictly less than 0.5, suitable to be the value of a conclusion in a countermodel; or the value 0.5 itself, not suitable to be the value of a premise or a conclusion in a countermodel.

Unlike usual designated-value or order-theoretic understandings of consequence on these models, understandings like Smith's that divide the values into *three* chunks, and require a countermodel to an argument to map its premises into one particular chunk and its conclusions to another, are not guaranteed to be *transitive* in any sense. Indeed, this is just the strategy exploited to produce nontransitivity in Frankowski (2004), Zardini (2008), Cobreros et al. (2012).⁸

In fact, there is a sense in which *every* consequence relation of a certain sort has a three-valued presentation: all that is required is that the consequence relation be *monotonic* (such that adding premises or conclusions can never make an invalid argument out of a valid one) and *reflexive* (such that every singleton set is a valid consequence of itself). Nothing like transitivity is required.⁹

⁷ Smith considers single-conclusion arguments only; this is the most natural generalization of his approach to a multiple-conclusion setting, which helps to bring out the symmetry in the definition.

⁸ See also Paris et al. (2009) for a definition of consequence relative to two thresholds in a probabilistic setting, to deal with inconsistent beliefs. The resulting consequence is also called "parameterized" there; our terminology below was introduced independently.

⁹ That *four* values suffice for every *monotonic* consequence relation, whether or not it is reflexive, is implicit in Proposition 2 of Humberstone (1988, p. 407); the move to three values to impose reflexivity is given (again, implicitly) later on the same page. See also Malinowski (2004), French and Ripley (2019), Blasio et al. (2017), Chemla and Egré (2019) for different presentations of this result. The representability of a monotonic logic by means of four values is connected to *Suszko's Thesis* (Suszko, 1977), which draws on a similar fact, imposing reflexivity and a strong form of transitivity to reduce the number of needed values to two. When only one of those two conditions is imposed, the number can be reduced to three (see Malinowski, 2009b, about dropping reflexivity, Frankowski, 2004, about dropping transitivity, and Tsuji, 1998; Blasio et al., 2017; French and Ripley, 2019; Chemla and Egré, 2019, for general results). Importantly, the fact that a consequence relation has *some* three-valued presentation does nothing to show that it has a *well-behaved* three-valued presentation. In particular, there is no requirement of compositionality; it might be that the value of a compound sentence floats completely free of the values of its components. In what follows, we do pursue some reductions to three-valued presentations, but these all maintain

As it happens, despite the presence of three inferential statuses, the possibility of nontransitivity is not realized in Smith (2008); Smith-consequence *is* transitive. In fact, it is precisely the usual consequence relation of classical logic (Smith, 2008, p. 222; this will also follow from our Theorem 10.1 below).

10.3.2 ST

The strict-tolerant approach to vagueness was first presented in Cobreros et al. (2012) as a(n implicitly) four-valued system.¹⁰ Here, we give a simpler three-valued formulation, following Cobreros et al. (2015a).

Definition 10.2 A model $\langle D, I \rangle$ is an *ST countermodel* to an argument $\Gamma \succ \Delta$ iff:

- $I(A) \in \{0, 0.5, 1\}$ for every formula A ;
- $I(\gamma) = 1$ for every $\gamma \in \Gamma$; and
- $I(\delta) = 0$ for every $\delta \in \Delta$.

If there is no ST countermodel to an argument $\Gamma \succ \Delta$, then the argument is *ST valid*, written $\Gamma \models_{ST} \Delta$. □

Note that the first clause of this definition amounts to throwing out all models that use any values other than those in $\{0, 0.5, 1\}$. (If such models cannot be countermodels, then there is no reason to attend to them at all when evaluating an argument for ST validity.) The remaining three-valued models are a familiar sort: they are *strong Kleene* models (see e.g. Beall & van Fraassen, 2003; Priest, 2008). To present ST in its own right, strong Kleene models are the simplest (model-theoretic) tool, but here we work within the broader space of models, to preserve convenient links with other consequence relations.

Note that every ST countermodel to an argument is also a Smith countermodel to that argument; it follows immediately that $\Gamma \models_S \Delta$ implies $\Gamma \models_{ST} \Delta$. As it happens, the converse holds as well; both consequence relations are exactly the familiar consequence relation of first-order classical logic. (Again, this will follow from Theorem 10.1.)

compositionality; the reductions we use are all truth-functional in the sense of Chemla and Egré (2019).

¹⁰ In Cobreros et al. (2012), there are three different satisfaction relations between models and formulas: tolerant, classical, and strict. Each is implied by the one(s) that follow it. As a result, there are four statuses a formula can have on a model: it can be strictly satisfied, classically but not strictly satisfied, tolerantly but not classically satisfied, or not satisfied at all. The status of each compound sentence is determined by the statuses of its components. This is the sense in which the system is four-valued.

10.4 Parameterised Consequence

In this section, we turn to a broad family of consequence relations, of which Smith-consequence and ST-consequence are two instances. A third instance will also be helpful:

Definition 10.3 A model $\langle D, I \rangle$ is a *classical countermodel* to an argument $\Gamma \succ \Delta$ iff:

- $I(A) \in \{0, 1\}$ for every formula A ;
- $I(\gamma) = 1$ for every $\gamma \in \Gamma$; and
- $I(\delta) = 0$ for every $\delta \in \Delta$.

If there is no classical countermodel to an argument $\Gamma \succ \Delta$, then the argument is *classically valid*, written $\Gamma \models_C \Delta$. \square

Inspection of our models reveals that this is just the usual notion of a classical countermodel, and so the usual notion of (first-order) classical validity.

To move to the general case, we need the notion of a set of values being *closed*:

Definition 10.4 A set $V \subseteq [0, 1]$ is *closed* iff it is closed under greatest lower bound, least upper bound, and the function $\neg(x) = 1 - x$. \square

The import of this definition is contained in the following fact:

Fact 10.1 If $V \subseteq [0, 1]$ is closed, then for any model $\langle D, I \rangle$: if for all n , for every n -ary predicate P , $I(P) \in V^{(D^n)}$, then for all formulas A , $I(A) \in V$. \square

Proof Induction on A 's formation. (Note that binary minimum and maximum are special cases of greatest lower bound and least upper bound, respectively.) \square

That is, for closed V , if a model assigns predicate values built only from values in V , then the entire model will assign values only from V . This notion is convenient for identifying usable selections from the value space $[0, 1]$. Note that every finite subset of $[0, 1]$ that is closed under the function $\neg(x) = 1 - x$ is closed, since every finite subset is closed under greatest lower bound and least upper bound (of sets). (In finite cases, these simply amount to minimum and maximum, respectively.) In particular, both $\{0, 0.5, 1\}$ and $\{0, 1\}$ are closed.

We consider a parameterised notion of countermodel. Our parameters have three coordinates: a set V of values; a set T of values for premises to take in a countermodel; and a set F of values for conclusions to take in a countermodel. Given these parameters, a countermodel is a model that takes values from V , maps all premises into T , and maps all conclusions into F . We will not consider all possible ways of doing this, but we will still consider quite a wide range.

Definition 10.5 A set $X \subseteq [0, 1]$ is an *upset* iff whenever $x \in X$ and $x < y \leq 1$, then $y \in X$; it is a *downset* iff whenever $x \in X$ and $x > y \geq 0$, then $y \in X$. \square

Definition 10.6 A *parameter* is a triple $\langle V, T, F \rangle$ such that:

- $\{0, 1\} \subseteq V \subseteq [0, 1]$ is closed;
- $1 \in T \subseteq (.5, 1]$ is an upset; and
- $0 \in F \subseteq [0, 0.5)$ is a downset.

□

Definition 10.7 Given a parameter $\mathcal{P} = \langle V, T, F \rangle$, a model $\langle D, I \rangle$ is a *\mathcal{P} countermodel* to an argument $\Gamma \succ \Delta$ iff:

- $I(A) \in V$ for every formula A ;
- $I(\gamma) \in T$ for every $\gamma \in \Gamma$; and
- $I(\delta) \in F$ for every $\delta \in \Delta$.

If there is no \mathcal{P} countermodel to an argument $\Gamma \succ \Delta$, then the argument is *\mathcal{P} valid*, written $\Gamma \models_{\mathcal{P}} \Delta$. □

Note that all three of our examples fit this mould: for Smith-consequence, the parameter is $\langle [0, 1], (0.5, 1], [0, 0.5) \rangle$; for ST-consequence, $\langle \{0, 0.5, 1\}, \{1\}, \{0\} \rangle$; and for classical consequence, $\langle \{0, 1\}, \{1\}, \{0\} \rangle$. By letting S, ST, and C simply *be* these parameters, we can see Definitions 10.1, 10.2, and 10.3 all as special cases of Definition 10.7.¹¹

We proceed to characterize \mathcal{P} validity for an arbitrary parameter, by way of a definition and a pair of lemmas.

Definition 10.8 A model $M' = \langle D, I' \rangle$ *crispifies* a model $M = \langle D, I \rangle$ (written $M \preccurlyeq M'$) iff for all n , for all n -ary predicates P :

- $I'(P) \in \{0, 1\}^{(D^n)}$;
- if $I(P)((d_1, \dots, d_n)) > 0.5$, then $I'(P)((d_1, \dots, d_n)) = 1$; and
- if $I(P)((d_1, \dots, d_n)) < 0.5$, then $I'(P)((d_1, \dots, d_n)) = 0$.

□

Note that when $M \preccurlyeq M'$, then M' assigns predicates only values from $\{0, 1\}$; since this set is closed, we already know that M' thus assigns every formula some value in $\{0, 1\}$. But something more interesting is happening here as well:

Lemma 10.1 *If $M \preccurlyeq M'$, then for every formula A , if $M(A) > 0.5$ then $M'(A) = 1$, and if $M(A) < 0.5$ then $M'(A) = 0$.* □

Proof Induction on A 's formation. □

Lemma 10.2 *Every model can be crispified. That is, for any M , there is some M' such that $M \preccurlyeq M'$.* □

¹¹ Our definition of a parameter rules out the overlap between T and F . By allowing overlap, and taking $T \subseteq (0, 1]$, and $F \subseteq [0, 1)$, we could retrieve another notion of entailment explored in the literature, namely TS entailment (see Cobreros et al., 2012), also known as Q-consequence (Malinowski, 1990, 2009a), then expressible as $\langle \{0, 0.5, 1\}, \{0.5, 1\}, \{0, 0.5\} \rangle$. We leave an exploration of TS entailment for another occasion.

Proof Definition 10.8 already tells us how to crispify each case where $I(P)$ assigns a value different from 0.5; for cases where $I(P)$ assigns exactly 0.5, either 1 or 0 will do. \square

We are now ready to characterize \mathcal{P} validity:

Theorem 10.1 *For any parameter \mathcal{P} , $\Gamma \models_{\mathcal{P}} \Delta$ iff $\Gamma \models_C \Delta$.* \square

Proof For each direction, we show the contrapositive.

LTR: suppose $\Gamma \not\models_C \Delta$. Then there is some classical countermodel M to $\Gamma \succ \Delta$, some M such that $M(A) \in \{0, 1\}$ for every formula A , $M(\gamma) = 1$ for every $\gamma \in \Gamma$ and $M(\delta) = 0$ for every $\delta \in \Delta$. But since $\mathcal{P} = \langle V, T, F \rangle$ is a parameter, this gives $M(A) \in V$ for every formula A (since $\{0, 1\} \subseteq V$), $M(\gamma) \in T$ for every $\gamma \in \Gamma$ (since $1 \in T$), and $M(\delta) \in F$ for every $\delta \in \Delta$ (since $0 \in F$). Thus, M is a \mathcal{P} countermodel to $\Gamma \succ \Delta$, and so $\Gamma \not\models_{\mathcal{P}} \Delta$.

RTL: suppose $\Gamma \not\models_{\mathcal{P}} \Delta$. Then there is some \mathcal{P} countermodel M to $\Gamma \succ \Delta$. This requires that $M(\gamma) \in T$ for every $\gamma \in \Gamma$ and $M(\delta) \in F$ for every $\delta \in \Delta$. Since \mathcal{P} is a parameter, this gives $M(\gamma) > 0.5$ for every $\gamma \in \Gamma$ and $M(\delta) < 0.5$ for every $\delta \in \Delta$. By Lemma 10.2, there is some model M' such that $M \preccurlyeq M'$. By Lemma 10.1, $M'(\gamma) = 1$ for every $\gamma \in \Gamma$ and $M'(\delta) = 0$ for every $\delta \in \Delta$. Since $M'(A) \in \{0, 1\}$ for every formula A , M' is a classical countermodel to $\Gamma \succ \Delta$, and so $\Gamma \not\models_C \Delta$. \square

From this perspective, it is no surprise that Smith-consequence and ST-consequence both turn out to be exactly classical logic; these are just two pinpricks of light shed on the broader phenomenon here, which is that *every* parameterised consequence relation is exactly classical logic.

This immediately yields an n -valued presentation of classical logic for every $n \geq 2$. Let $V_n = \{0, 1/(n-1), \dots, (n-2)/(n-1), 1\}$. Then V_n has n members, and $\mathcal{P} = \langle V_n, \{1\}, \{0\} \rangle$ is a parameter. It also reveals that Smith's choice of the parameter $\langle \{0, 1\}, (0.5, 1], [0, 0.5) \rangle$ is logically arbitrary, even given his choice of the value space $[0, 1]$. Many choices for the parameter's last two coordinates would have yielded the same consequence relation.

Smith offers more motivation for his choice than simply the consequence relation it yields. He understands a sentence with value ≥ 0.5 as 'assertion grade' (fit to assert) and a sentence with value > 0.5 as 'inference grade' (fit to infer from); the idea is that a valid argument whose premises are all inference grade must have some conclusion that is assertion grade. Why these particular values?

The advantage of my proposal... is that it is *minimal*. If a sentence S is at least 0.5 true, then one cannot make a truer statement by asserting the negation of S than by asserting S . What more than this could be required for a statement to be 'assertion grade'...? Any higher standard would need further justification, and I cannot see what such justification would consist in. Now, given that we have set the cut-off for assertion grade statements at 0.5, and want to make the cut-off for inference grade statements strictly *higher* than this, the *minimal* cut-off for inference grade statements will be the one I have proposed: they must be *more than* 0.5 true. Again, any higher standard would need further justification, and I cannot see what such justification would consist in. (Smith, 2008, p. 224, emphases in original)

We don't see, however, that minimality in this sense is any advantage to a proposal at all.¹²

We conclude, then, that for capturing the logical behaviour common to ST and to Smith's approach before we take account of similarity predicates, any parameter will do as well as any other. All yield the same logic (ordinary classical logic), and we see no other potential reason to choose between them (so long as the values are interpreted in a way that circumvents the jolt problem, as discussed in Sect. 10.2).

10.4.1 Tolerant Logics

Here, we move on to consider the logic of P -similarity, registering the connections between \sim_P and P . For tolerant logics, we impose a connection between the values assigned to $t \sim_P u$, to Pt , and to Pu , for every predicate P and terms t, u . It is perhaps not immediately apparent how to understand these connections with regard to Smith-consequence, so we pursue an indirect approach. First we consider the situation as developed in Cobreros et al. (2012, 2015a). Then we turn to the general case, first looking at parameterised consequence relations in their full generality, and then narrowing in on a particular class of them; ST-consequence will be seen as a member of this wider class.

We begin by narrowing our space of models to ensure that \sim relations are one and all reflexive and symmetric, in the following sense: for all predicates P and all terms t, u : $I(t \sim_P t) = 1$ and $I(t \sim_P u) = I(u \sim_P t)$. From here forward, we ignore models that do not obey these restrictions.¹³

10.4.1.1 ST \sim

Definition 10.9 A model M obeys the $S\sim$ restriction iff for all predicates P and terms t, u : if $M(Pt) = 1$ and $M(Pu) = 0$, then $M(t \sim_P u) = 0$. \square

Intuitively, the ST \sim restriction ensures that if t and u are so P -unlike as to go all the way from 1 to 0 in their P -value, then they must not be at all P -similar.

¹² The above-quoted passage is the full discussion of the issue in Smith (2008), except for (p. 250, fn. 57, emphasis in original): '[Earlier], I said that a sentence is 'assertion grade' ... if its degree of truth is greater than or equal to 0.5. This does *not* mean that if a sentence S has a degree of truth of 0.5 or greater, then an ... assertion of S is acceptable. Rather, the idea is that a sentence is 'assertion grade' ... if the level of confidence appropriate in an utterance of the sentence is at least as high as the level of confidence appropriate in an utterance of its negation.' We set aside further discussion of how exactly to understand 'assertion grade' and 'inference grade' on Smith's account.

¹³ This of course already results in more valid arguments than we already had. Look ahead to Fig. 10.1. By removing the rule Tol from the sequent calculus there, you arrive at a calculus sound and complete for every \mathcal{P} consequence relation obeying these \sim restrictions. (Yes, they are all the same; the argument is the same as for Theorem 10.1, mutatis mutandis).

Definition 10.10 A model M is an ST_{\sim} *countermodel* to an argument $\Gamma \succ \Delta$ iff:

- M obeys the ST_{\sim} restriction, and
- M is an ST counterexample to $\Gamma \succ \Delta$.

If there is no ST_{\sim} countermodel to an argument $\Gamma \succ \Delta$, then the argument is ST_{\sim} *valid*, written $\Gamma \models_{ST_{\sim}} \Delta$. \square

It will be convenient later to have a proof system for ST_{\sim} . We will work with a sequent calculus, given in Fig. 10.1. In the figure, t and u can be any terms, and a must be an *eigenvariable*: a variable that does not occur free in the conclusion-sequent of the rule.

Fact 10.2 The sequent calculus in Fig. 10.1 is sound and complete for ST_{\sim} . \square

Proof Both facts are straightforward to show in the usual ways. In particular, completeness can be proved following the method of Takeuti (1987), which builds a countermodel from an underivable sequent. For the connectives and quantifiers, this method is adapted to strong Kleene models in Ripley (2013a).

The only needed addition is to show that the resulting model handles \sim appropriately, but this is ensured by the \sim -involving rules. In particular, \sim ref ensures that the resulting model assigns 1 to all formulas of the form $t \sim_P t$, the \sim -sym rules ensure that the resulting model assigns the same value to $t \sim_P u$ as to $u \sim_P t$, and the Tol rule guarantees that the resulting model obeys the ST_{\sim} restriction. \square

ST_{\sim} is an *expansion* of first-order classical logic. ST_{\sim} countermodels are all ST countermodels, so ST validity (which we know is classical validity from Theorem 10.1) implies ST_{\sim} validity. ST_{\sim} also has the nice feature that it validates the principle of tolerance. This is so whether we consider instances of tolerance as *arguments* (since $Pa, a \sim_P b \models_{ST_{\sim}} Pb$) or as *quantified conditionals* (since $\models_{ST_{\sim}} \forall x \forall y ((Px \wedge x \sim_P y) \rightarrow Py)$). The main form of tolerance we're interested in, though, is the *metainferential* form, expressed by the rule Tol in Fig. 10.1.

Definition 10.11 A consequence relation is *tolerant* iff it is closed under the rule Tol. \square

Since Tol is part of a sound and complete sequent calculus for ST_{\sim} , ST_{\sim} is tolerant. In this calculus, the other forms of tolerance follow from the metainferential form, via the rules Id, \wedge L, \rightarrow R, and \forall R. We will return to this in Sect. 10.4.3.

There is a natural worry at this point: doesn't ST_{\sim} fall right into the sorites paradox? After all, quantified conditional tolerance plus classical logic is typically thought to be enough for trouble, yet ST_{\sim} validates these. But ST_{\sim} has an escape route: the structure of its countermodels allows it to be *nontransitive*. Indeed, ST_{\sim} is nontransitive, and this allows it to escape the looming trouble. We do not pursue the ups and downs of nontransitivity here; Cobreros et al. (2012, 2015a) offer related discussion.

Structural rules:

$$\text{Id: } \frac{}{A \succ A} \quad \text{K: } \frac{\Gamma \succ \Delta}{\Gamma, \Gamma' \succ \Delta, \Delta'}$$

Operational rules:

$$\neg\text{L: } \frac{\Gamma \succ A, \Delta}{\Gamma, \neg A \succ \Delta} \quad \neg\text{R: } \frac{\Gamma, A \succ \Delta}{\Gamma \succ \neg A, \Delta}$$

$$\wedge\text{L: } \frac{\Gamma, A, B \succ \Delta}{\Gamma, A \wedge B \succ \Delta} \quad \wedge\text{R: } \frac{\Gamma \succ A, \Delta \quad \Gamma \succ B, \Delta}{\Gamma \succ A \wedge B, \Delta}$$

$$\vee\text{L: } \frac{\Gamma, A \succ \Delta \quad \Gamma, B \succ \Delta}{\Gamma, A \vee B \succ \Delta} \quad \vee\text{R: } \frac{\Gamma \succ A, B, \Delta}{\Gamma \succ A \vee B, \Delta}$$

$$\rightarrow\text{L: } \frac{\Gamma \succ A, \Delta \quad \Gamma, B \succ \Delta}{\Gamma, A \rightarrow B \succ \Delta} \quad \rightarrow\text{R: } \frac{\Gamma, A \succ B, \Delta}{\Gamma \succ A \rightarrow B, \Delta}$$

$$\forall\text{L: } \frac{\Gamma, A(t) \succ \Delta}{\Gamma, \forall x A(x) \succ \Delta} \quad \forall\text{R: } \frac{\Gamma \succ A(a), \Delta}{\Gamma \succ \forall x A(x), \Delta}$$

$$\exists\text{L: } \frac{\Gamma, A(a) \succ \Delta}{\Gamma, \exists x A(x) \succ \Delta} \quad \exists\text{R: } \frac{\Gamma \succ A(t), \Delta}{\Gamma \succ \exists x A(x), \Delta}$$

Similarity rules:

$$\sim\text{ref: } \frac{\Gamma, t \sim_P t \succ \Delta}{\Gamma \succ \Delta}$$

$$\sim\text{symL: } \frac{\Gamma, t \sim_P u \succ \Delta}{\Gamma, u \sim_P t \succ \Delta} \quad \sim\text{symR: } \frac{\Gamma \succ t \sim_P u, \Delta}{\Gamma \succ u \sim_P t, \Delta}$$

$$\text{Tol: } \frac{\Gamma \succ t \sim_P u, \Delta}{\Gamma, Pt \succ Pu, \Delta}$$

Fig. 10.1 A sequent calculus for ST~

10.4.1.2 Parameter Tolerance

Our concern here is to explore various avenues for restricting parameterised consequence so as to respect the connection between P and \sim_P . We take as our paradigm the ST_{\sim} restriction, and generalise it to what we call *parameter tolerance*.

In any parameter $\mathcal{P} = \langle V, T, F \rangle$, T is some sort of positive status, and F some kind of negative status. The force of \mathcal{P} validity is to guarantee that if the premises all have the positive status, then some conclusion lacks the negative status. Parameter tolerance is parasitic on these statuses: it guarantees that if Pt has the positive status while Pu has the negative one, then $t \sim_P u$ must also have the negative one.

Definition 10.12 For a parameter $\mathcal{P} = \langle V, T, F \rangle$, a model M is \mathcal{P} *tolerant* iff for all predicates P , for all terms t, u : if $M(Pt) \in T$ and $M(Pu) \in F$, then $M(t \sim_P u) \in F$. \square

We can use parameter tolerance to get a new range of parameterised consequence relations:

Definition 10.13 A model M is a \mathcal{P} *tolerant countermodel* to an argument $\Gamma \succ \Delta$ iff:

- M is \mathcal{P} tolerant, and
- M is a \mathcal{P} countermodel to $\Gamma \succ \Delta$.

If there is no \mathcal{P} tolerant countermodel to an argument $\Gamma \succ \Delta$, then the argument is \mathcal{P} *tolerant valid*, written $\Gamma \models_{\mathcal{P}} \Delta$. \square

Note that with ST understood as the parameter $\langle \{0, 0.5, 1\}, \{1\}, \{0\} \rangle$, as before, Definition 10.13 subsumes Definition 10.10. Parameter tolerance brings with it metainferential tolerance.

Fact 10.3 For every parameter \mathcal{P} , $\models_{\mathcal{P}}$ is tolerant. \square

Proof Let $\mathcal{P} = \langle V, T, F \rangle$, and suppose that $\Gamma \cup \{Pt\} \not\models_{\mathcal{P}} \{Pu\} \cup \Delta$. Then there is a \mathcal{P} tolerant model M such that $M(\gamma) \in T$ for every $\gamma \in \Gamma \cup \{Pt\}$ and $M(\delta) \in F$ for every $\delta \in \Delta \cup \{Pu\}$. In particular, $M(Pt) \in T$ and $M(Pu) \in F$. Since M is \mathcal{P} tolerant, then, $M(t \sim_P u) \in F$. But then $\Gamma \not\models_{\mathcal{P}} t \sim_P u, \Delta$, since M is a countermodel to this argument as well. \square

As the proof reveals, \mathcal{P} tolerance is just what is needed to guarantee metainferential tolerance. Since all parameterised consequence relations are reflexive, this is enough to guarantee argument-form tolerance as well. (For other forms of tolerance, we need more restrictions yet; leave those to one side for now.)

Moreover, these consequence relations are tied quite tightly to classical logic, as Fact 10.4 records:

Fact 10.4 For every parameter \mathcal{P} , $\models_{\mathcal{P}}$ is a *conservative extension* of classical logic: if $\Gamma \models_{\mathcal{P}} \Delta$ and $\Gamma \not\models_C \Delta$, then some \sim predicate occurs in $\Gamma \cup \Delta$. \square

Proof Suppose that $\Gamma \cup \Delta$ contains no occurrences of any \sim_P relation, and suppose that $\Gamma \not\models_C \Delta$. Then there is a classical countermodel M for the language without \sim relations. Extend M to a model M' of the full language as follows: for all terms t, u and predicates P , if $M(t) = M(u)$, then $M'(t \sim_P u) = 1$; else $M(t \sim_P u) = 0$. M' is \mathcal{P} tolerant, as is quick to check. But $M'(A) = M(A)$ for all $A \in \Gamma \cup \Delta$. So M' is a \mathcal{P} tolerant countermodel: $\Gamma \not\models_{\mathcal{P}} \Delta$. \square

The \mathcal{P} tolerant consequence relations, then, are a way of extending classical logic to take account of the special behaviour of \sim relations, without messing with anything that doesn't involve these relations.

10.4.2 Sorites

As we flagged above, ST_{\sim} handles the combination of classical logic and tolerance without running into sorites problems via *nontransitivity*. A natural question that arises at this point, then, is: which of the $\models_{\mathcal{P}}^{\sim}$ relations work the same way? The answer is: the *proper* ones.

Definition 10.14 A parameter $\mathcal{P} = \langle V, T, F \rangle$ is *proper* iff $V \not\subseteq T \cup F$. The consequence relation $\models_{\mathcal{P}}^{\sim}$ is *proper* iff \mathcal{P} is. \square

Proper parameters are parameters with some value that is neither in T nor F ; such a value is needed for any counterexample to transitivity to arise. With this notion in hand, we have the following containment result:

Fact 10.5 For every proper parameter \mathcal{P} , $\models_{\mathcal{P}}^{\sim} \subseteq \models_{ST}^{\sim}$. \square

Proof Let $\mathcal{P} = \langle V, T, F \rangle$, and suppose $\Gamma \not\models_{ST}^{\sim} \Delta$. Then there is a model M meeting the ST_{\sim} restriction such that: $M(\gamma) = 1$ for all $\gamma \in \Gamma$, and $M(\delta) = 0$ for all $\delta \in \Delta$. Since $V \not\subseteq T \cup F$, there must be some value in V that is in neither T nor F ; call it x . Since V is closed, $1 - x \in V$ as well. Now consider the model M' defined as follows: for all n for all n -ary predicates P ,

- If $M(P)(\langle d_1, \dots, d_n \rangle) = 1$, then $M'(P)(\langle d_1, \dots, d_n \rangle) = 1$;
- if $M(P)(\langle d_1, \dots, d_n \rangle) = 0$, then $M'(P)(\langle d_1, \dots, d_n \rangle) = 0$; and
- if $M(P)(\langle d_1, \dots, d_n \rangle) = 0.5$, then $M'(P)(\langle d_1, \dots, d_n \rangle) = x$.

Note that $\{0, x, 1 - x, 1\}$ is closed, so M' assigns only these values. Moreover, we can show by induction on A 's formation that for all formulas A , if $M(A) = 1$ or 0 , then $M'(A) = M(A)$. Thus, M' still assigns 1 to everything in Γ and 0 to everything in Δ . Moreover, it assigns values only from V , so it is a \mathcal{P} countermodel. It remains only to show that M' is \mathcal{P} tolerant. Suppose $M'(Pt) \in T$ and $M'(Pu) \in F$. There are three cases:

- If $M(Pt) = 1$ and $M(Pu) = 0$, we have $M(t \sim_P u) = 0$, since M meets the ST_{\sim} restriction, and so $M'(t \sim_P u) = 0 \in F$.

- If $M(Pt) \neq 1$, then $M(Pt) = 0.5$ or 0 , and so $M'(Pt) = x$ or 0 . But either way, $M'(Pt) \notin T$; contradiction.
- If $M(Pu) \neq 0$, then $M(Pu) = 0.5$ or 1 , and so $M'(Pu) = x$ or 1 . But either way, $M'(Pu) \notin F$; contradiction.

M' is thus a \mathcal{P} tolerant countermodel to $\Gamma \succ \Delta$, and so $\Gamma \not\models_{\mathcal{P}} \Delta$. \square

The sequent calculus in Fig. 10.1 is thus complete for every proper $\models_{\mathcal{P}}$. (It is not, however, sound for all of them; we return to this in Sect. 10.4.3.)

Fact 10.5 does not hold for improper parameters. This is because improper parameters walk straight into the business end of the sorites paradox.

Fact 10.6 For any improper \mathcal{P} and any n terms t_1, \dots, t_n : $Pt_1, t_1 \sim_P t_2, t_2 \sim_P t_3, \dots, t_{n-1} \sim_P t_n \models_{\mathcal{P}} Pt_n$. \square

Proof Let $\mathcal{P} = \langle V, T, F \rangle$, and suppose the claim fails. Then there is some \mathcal{P} tolerant countermodel M . It must be that $M(Pt_1) \in T$ and that $M(Pt_n) \in F$, hence $M(Pt_n) \notin T$; so there is a first $i \leq n$ such that $M(Pt_i) \notin T$. Since i is the first such, $M(Pt_{i-1}) \in T$, and since \mathcal{P} is improper, $M(Pt_i) \in F$. Since M is \mathcal{P} tolerant, then, $M(t_{i-1} \sim_P t_i) \in F$. But then M cannot be a countermodel to this argument, since $t_{i-1} \sim_P t_i$ is among its premises. Contradiction. \square

So while improper parameters give rise to tolerant conservative extensions of classical logic, they are not useful for exploring sorites sequences, as they fall victim to the sorites paradox, by requiring that if the first member of a sorites sequence for P is P , then so is the last. Much of what follows holds for parameters whether or not they are proper, so we do not restrict our attention only to proper parameters, but we do think that improper parameters are unlikely to be of any help in treating vagueness, since they do not avoid the key problem posed by sorites sequences.

On the other hand, proper parameters give rise to nontransitivity just where it is needed. Consider the form of transitivity embodied in the rule of *Cut*:

$$\text{Cut: } \frac{\Gamma \vdash A, \Delta \quad \Gamma', A \vdash \Delta'}{\Gamma, \Gamma' \vdash \Delta, \Delta'}$$

Fact 10.7 For any proper \mathcal{P} , $\models_{\mathcal{P}}$ is not closed under *Cut*.

Proof Since every $\models_{\mathcal{P}}$ is reflexive and tolerant, they all validate the arguments $Pt_1, t_1 \sim_P t_2 \succ Pt_2$ and $Pt_2, t_2 \sim_P t_3 \succ Pt_3$. If we could apply *Cut* to these, we would reach $Pt_1, t_1 \sim_P t_2, t_2 \sim_P t_3 \succ Pt_3$. But this last is not ST_{\sim} valid, so not \mathcal{P} tolerant valid for any proper \mathcal{P} by Fact 10.5. \square

Proper $\models_{\mathcal{P}}$ relations thus handle sorites reasoning just like \models_{ST} : by validating each step of the reasoning, but refusing to allow them to be chained together via *Cut*.

10.4.3 Metainferences

Theorem 10.1 guarantees that every \mathcal{P} tolerant consequence relation is *argument classical*: they all validate every instance of every classically-valid argument. But this is just one way to be classical; we might (and probably should) want more.

As we pointed out above, the sequent calculus in Fig. 10.1 is complete for all proper parameter tolerant consequence, but it is not *sound* for all of them. This may at first seem counterintuitive, since every $\models_{\mathcal{P}}^{\sim}$ strengthens classical logic and obeys the similarity rules, while these rules are the only rules in Fig. 10.1 *not* sound for classical logic. But counterintuitive or not, it is the case; we pause here to sort this out.

We can use the calculus of Fig. 10.1 to identify certain *metainferences* important to classicality. (A *metainference* is a property that a consequence relation may or may not be closed under. For discussion, see e.g. Field, 2008, Scharp, 2013, Cobreros et al., 2013; Barrio et al., 2015).

Fact 10.8 Every \mathcal{P} tolerant consequence relation is closed under the metainferences given by the rules Id, K, \sim ref, \sim symL, and \sim symR. \square

Proof Straightforward. (For the \sim rules, recall that we have restricted our models so that $M(t \sim_P t) = 1$ and $M(t \sim_P u) = M(u \sim_P t)$.) \square

Facts 10.3 and 10.8 ensure that every \mathcal{P} tolerant consequence relation obeys the structural rules and similarity rules given in Fig. 10.1. This leaves the operational rules; these determine what we will call the *operational metainferences*. We say that a consequence relation is *operationally classical* iff it is closed under all of these operational metainferences.

It is entirely possible for a consequence relation to be argument classical *without* being operationally classical. For example, consider the metainference determined by the rule \rightarrow R, and consider the parameter $\mathcal{P} = \langle [0, 1], \{1\}, [0, 0.5] \rangle$. We have $Pa \wedge a \sim_P b \models_{\mathcal{P}} Pb$, but $\not\models_{\mathcal{P}} (Pa \wedge a \sim_P b) \rightarrow Pb$; for a countermodel to the latter, let $M(Pa) = M(a \sim_P b) = 0.6$ and $M(Pb) = 0.4$. Note that this model can be \mathcal{P} tolerant, since $M(Pa) \notin T$.

Here, then, we describe the situation for the operational metainferences, identifying sufficient conditions for a parameter \mathcal{P} to yield an operationally classical \mathcal{P} tolerant consequence relation.

Fact 10.9 Every parameterised consequence is closed under \wedge L, \wedge R, \vee L, \vee R, \forall L, and \exists R. \square

Proof Let $\mathcal{P} = \langle V, T, F \rangle$.

For \wedge L: A countermodel to the conclusion-sequent must assign some value in T to $A \wedge B$; but since this value is the minimum of the values of A and B , and since T is an upset, this model must assign some value in T to each of A and B , and so be a countermodel to the premise-sequent.

For \wedge R: A countermodel to the conclusion-sequent must assign some value in F to $A \wedge B$; but since this value is the minimum of the values of A and B , this model

must assign that value (and so some value in F) to at least one of A or B , and so be a countermodel to at least one premise-sequent.

For $\forall L$: A countermodel M to the conclusion-sequent must assign some value in T to $\forall x A(x)$; but since this value is a lower bound for the values of $A(x)$ in x -variants of M , and since T is an upset, all of these x -variants must assign some value in T to $A(x)$. There must be some x -variant M' such that $M'(x) = M(t)$, so $M(A(t)) \in T$; M is thus a countermodel to the premise-sequent.

$\forall L$ is similar to $\wedge R$; $\vee R$ to $\wedge L$; and $\exists R$ to $\forall L$. \square

This leaves two kinds of metainferences: the *negative* ones $\neg L$, $\neg R$, $\rightarrow L$, and $\rightarrow R$; and the *eigenvariable* ones $\forall R$ and $\exists L$. Parameterised tolerant consequence relations as such are not guaranteed to be closed under any of these. The trouble with the negative ones is that there is no connection between $x \in T$ and $1 - x \in F$; and the trouble with the eigenvariable ones is that the lub or glb of a set X might be in T or F without any member of X being so. To ensure that the remaining operational metainferences work, we need to tighten up our parameters in these two ways.

Definition 10.15 A parameter $\mathcal{P} = \langle V, T, F \rangle$ is *symmetric* iff for all $x \in [0, 1]$: $x \in T$ iff $1 - x \in F$. The consequence relation $\models_{\mathcal{P}}$ is *symmetric* iff \mathcal{P} is. \square

In a symmetric parameter $\langle V, T, F \rangle$, T and F are mirror images of each other, with the function $\neg(x) = 1 - x$ serving as the mirror.

Fact 10.10 \mathcal{P} is symmetric iff $\models_{\mathcal{P}}$ is closed under $\neg L$ and $\neg R$ iff $\models_{\mathcal{P}}$ is closed under $\rightarrow L$ and $\rightarrow R$. \square

Proof LTR: Let $\mathcal{P} = \langle V, T, F \rangle$. For $\neg L$: A countermodel M to the conclusion-sequent must have $M(\neg A) \in T$. Since it is symmetric, this gives $1 - M(\neg A) \in F$. But $M(A) = 1 - M(\neg A)$, so $M(A) \in F$, and M is a countermodel to the premise-sequent.

$\neg R$ is similar to $\neg L$.

Since $M(A \rightarrow B) = M(\neg A \vee B)$ in every model, it suffices for $\rightarrow L$ and $\rightarrow R$ to show that $\neg A \vee B$ obeys these rules. But this can be derived from $\neg L$, $\neg R$, $\vee L$, and $\vee R$ as follows:

$$\begin{array}{c} \neg L: \frac{\Gamma \vdash A, \Delta}{\Gamma, \neg A \vdash \Delta} \\ \neg L: \frac{\Gamma, \neg A \vdash \Delta \quad \Gamma, B \vdash \Delta}{\Gamma, \neg A \vee B \vdash \Delta} \end{array} \quad \begin{array}{c} \neg R: \frac{\Gamma, A \vdash B, \Delta}{\Gamma \vdash \neg A, B, \Delta} \\ \vee R: \frac{\Gamma \vdash \neg A, B, \Delta}{\Gamma \vdash \neg A \vee B, \Delta} \end{array}$$

RTL: By Fact 10.8, we know that $\models_{\mathcal{P}}$ obeys Id and Tol for any \mathcal{P} . If $\models_{\mathcal{P}}$ is also closed under $\neg L$, it follows that $Pa, a \sim_{\mathcal{P}} b, \neg Pb \models_{\mathcal{P}}$, so providing a counterexample to this would show that $\models_{\mathcal{P}}$ isn't closed under $\neg L$. Similarly, if $\models_{\mathcal{P}}$ is closed under $\neg R$, it follows that $a \sim_{\mathcal{P}} b \models_{\mathcal{P}} Pb, \neg Pa$, so providing a counterexample to this would show that $\models_{\mathcal{P}}$ isn't closed under $\neg R$.

So suppose $\mathcal{P} = \langle V, T, F \rangle$ is not symmetric. Then either there is $x \in T$ but $1 - x \notin F$, or there is $x \in F$ but $1 - x \notin T$. Take a model M with its domain

containing just the terms a and b ; let the terms a and b denote themselves, and let all other constant terms denote a .

If there is $x \in T$ but $1 - x \notin F$, then let $M(P)(a) = 1$, $M(P)(b) = 1 - x$, and $M(\sim_P)(a, b) = M(\sim_P)(b, a) = 1$. For all other predicates, let them take everything to 1. This clearly gives a \mathcal{P} counterexample to $Pa, a \sim_P b, \neg Pb \succ$; it is \mathcal{P} tolerant as well, since the only real risk to tolerance comes from $a \sim_P b$, and while $M(Pa) \in T$, we do not have $M(Pb) \in F$. So in this case $\models_{\mathcal{P}}$ is not closed under $\neg L$.

Or if there is $y \in F$ but $1 - y \notin T$, then let $M(P)(a) = y$, $M(P)(b) = 0$, and $M(\sim_P)(a, b) = M(\sim_P)(b, a) = 1$. For all other predicates, let them take everything to 1. This clearly gives a \mathcal{P} counterexample to $a \sim_Q b \succ Pb, \neg Pa$; it is \mathcal{P} tolerant as well, since the only real risk to tolerance comes from $a \sim_Q b$, and while $M(Pb) \in F$, we do not have $M(Pa) \in T$. So in this case $\models_{\mathcal{P}}$ is not closed under $\neg R$.

Similar arguments will show that in the first case $\models_{\mathcal{P}}$ is not closed under $\rightarrow L$ and in the second not under $\rightarrow R$. (We have $M(\neg(a \sim_P a)) = 0$, so $A \rightarrow \neg(a \sim_P a)$ is equivalent to $\neg A$ on any model.) \square

The negative metainferences addressed, only the eigenvariable metainferences remain.

Definition 10.16 A parameter $\mathcal{P} = \langle V, T, F \rangle$ is *open* iff for all $X \subseteq V$, if $\text{glb } X \in F$ then $F \cap X \neq \emptyset$, and if $\text{lub } X \in T$ then $T \cap X \neq \emptyset$.¹⁴ The consequence relation $\models_{\mathcal{P}}$ is *open* iff \mathcal{P} is. \square

Note that every parameter of the form $\langle [0, 1], (x, 1], [0, y) \rangle$ is open, as is every parameter $\langle V, T, F \rangle$ with finite V . But not all parameters are open; take $\langle [0, 1], [0.6, 1], [0, 0.4] \rangle$. Let $X = \{x \in [0, 1] : x < 0.6\}$. Then $\text{lub } X = 0.6 \in T$, but $T \cap X = \emptyset$.

Fact 10.11 \mathcal{P} is open iff $\models_{\mathcal{P}}$ is closed under $\forall R$ and $\exists L$. \square

Proof LTR: Let $\mathcal{P} = \langle V, T, F \rangle$. Suppose $\Gamma \not\models_{\mathcal{P}} \forall x A(x), \Delta$. Then there is some model M with: $M(\gamma) \in T$ for all $\gamma \in \Gamma$; $M(\delta) \in F$ for all $\delta \in \Delta$; and $M(\forall x A(x)) \in F$. Since $M(\forall x A(x)) = \text{glb}(\{M'(A(x)) : M' \text{ is an } x\text{-variant of } M\})$ and \mathcal{P} is open, there must be some x -variant M' of M such that $M'(A(x)) \in F$. Consider now the a -variant M'' of M such that $M''(a) = M'(x)$. Since a is an eigenvariable, M'' matches M on everything in Γ and Δ . So $\Gamma \not\models_{\mathcal{P}} A(x), \Delta$, since M'' is a \mathcal{P} tolerant countermodel.

$\exists L$ is similar.

RTL: By Fact 10.8, we know that $\models_{\mathcal{P}}$ obeys Id and Tol for any \mathcal{P} . If $\models_{\mathcal{P}}$ is also closed under $\exists L$, it follows that $\exists y Py, \forall x(x \sim_P a) \models_{\mathcal{P}} Pa$, so providing a counterexample to this would show that $\models_{\mathcal{P}}$ isn't closed under $\exists L$. Similarly, if $\models_{\mathcal{P}}$

¹⁴ This condition on T is that it be *Scott open* (see e.g. Vickers, 1989, p. 95), and this condition on F is the order-dual.

is closed under $\forall R$, it follows that $Pa, \forall x(a \sim_P x) \models_P \forall yPy$, so providing a counterexample to this would show that \models_P isn't closed under $\forall R$.

So suppose $\mathcal{P} = \langle V, T, F \rangle$ is not open. Then there is $X \subseteq V$ with either: $\text{lub } X \in T$ but $X \cap T = \emptyset$, or $\text{glb } X \in F$ but $X \cap F = \emptyset$. Take a model M with domain $X \cup \{a\}$, for some $a \notin X$; let every constant term of the language denote a . For every $x \in X$, let $M(P)(x) = x$, let $M(\sim_P)$ take every pair of objects to 1, and let all other predicates take every object to 1.

If $\text{lub } X \in T$ but $X \cap T = \emptyset$, then let $M(P)(a) = 0$. In this case, M is a \mathcal{P} counterexample to $\exists yPy, \forall x(x \sim_P a) \succ Pa$; it is \mathcal{P} tolerant since there is no z in the domain with $M(P)(z) \in T$.

If $\text{glb } X \in F$ but $X \cap F = \emptyset$, then let $M(P)(a) = 1$. In this case, M is a \mathcal{P} counterexample to $Pa, \forall x(a \sim_P x) \succ \forall yPy$; it is \mathcal{P} tolerant since there is no z in the domain with $M(P)(z) \in F$. \square

We now have a lot of pieces scattered around. Putting them together:

Theorem 10.2 *For any proper symmetric open parameter \mathcal{P} , $\models_{\mathcal{P}} = \models_{ST}^{\sim}$.* \square

Proof Fact 10.5 gives us that $\Gamma \models_{\mathcal{P}}^{\sim} \Delta$ implies $\Gamma \models_{ST}^{\sim} \Delta$ for proper \mathcal{P} ; it remains only to show the converse. Facts 10.3, 10.8, 10.9, 10.10, and 10.11 together establish that the sequent calculus in Fig. 10.1 is sound for symmetric open $\models_{\mathcal{P}}^{\sim}$. But since this calculus is complete for \models_{ST}^{\sim} (Fact 10.2), we're done. \square

That is, there is only one proper symmetric open \mathcal{P} tolerant consequence relation, and it is exactly the consequence relation of ST^{\sim} ! The consequence relation of ST^{\sim} thus occupies a natural place among parameterised consequence relations. The steps needed to dodge sorites trouble (properness) and ensure operational classicality (symmetry and openness) also ensure that the consequence relation of ST^{\sim} is the only choice. As these are natural desiderata, ST^{\sim} looms large.

Earlier, we identified Smith's approach as the parameterised consequence relation with the parameter $S = (\{0, 1\}, \{0.5, 1\}, \{0, 0.5\})$. Turning to its tolerant extension, we can see that this parameter is proper, symmetric, and open, so \models_S^{\sim} , like all such consequence relations, is the same as \models_{ST}^{\sim} . We thus reckon that anyone interested in taking a Smith-style approach to consequence in a language including \sim relations should arrive at ST^{\sim} as their desired consequence relation.

10.5 How Many Truth Values?

We have, in effect, presented many different model theories for the same consequence relation. Recall our earlier n -valued model theory for classical logic, based on the parameter $\langle V_n, \{1\}, \{0\} \rangle$ with $V_n = \{0, 1/(n-1), \dots, (n-2)/(n-1), 1\}$. This parameter is proper when $n \geq 3$, symmetric, and since V_n is finite, it is also open. So we also have n -valued presentations of ST^{\sim} for all $n \geq 3$. Of course, we also have continuum-valued presentations, provided by parameters of the form $\langle [0, 1], (x, 1), [0, 1-x] \rangle$, for $x \in [0.5, 1]$; these are all proper, symmetric, and

open, and Smith-consequence is among them, with $x = 0.5$. We can have countably-valued presentations as well. For example, let $V = \{x \in [0, 1] : \exists n \in \mathbb{N} \text{ s.t. } x = \frac{1}{2^n} \text{ or } x = \frac{2^n-1}{2^n}\}$. That is, $V = \{0, \dots, \frac{1}{16}, \frac{1}{8}, \frac{1}{4}, \frac{1}{2}, \frac{3}{4}, \frac{7}{8}, \frac{15}{16}, \dots, 1\}$. V is closed and countable. Now consider the parameter $\mathcal{P} = \langle V, [\frac{3}{4}, 1], [0, \frac{1}{4}] \rangle$. This parameter is proper, symmetric, and open, and so subject to Theorem 10.2.¹⁵

What to make of this situation? We argued in Sect. 10.2 that three-valued approaches, especially those that validate tolerance, need not face the jolt problem, if appropriately interpreted. We showed in Sect. 10.4 that ST_{\sim} is a three-valued approach that validates tolerance, and that it also captures exactly the arguments valid in Smith's continuum-valued framework, when it is extended with similarity predicates. There are two purposes, then, for which Smith's continuum-many values are otiose: they are not needed to avoid the jolt problem, nor are they needed to do any logical work. For each of these purposes, three values suffice.

Earlier, we distinguished between the referential and the inferential roles that truth values play in a theory of meaning. Cast in terms of that distinction, we see that Smith's theory uses infinitely many values as references of sentences, but our argument shows that only three values are needed at the inferential level to capture his consequence relation. By contrast, the strict-tolerant theory basically uses three values at both levels, referential and inferential: on our approach, the third value is assigned to a vague predicate to capture borderline status, and the same third value is used to validate the tolerance principle. Smith argues that any finite assignment of values to vague sentences will involve 'jolts' between distinct semantic values, but our contention is that a jolt-free approach is achieved first and foremost at the inferential level in the form of tolerance principles. In brief, the strict-tolerant account trims down Smith's referential apparatus to the minimum number of truth values needed at the inferential level.

This is not yet to say that three values will always be enough to satisfy any further desiderata you might consider for a theory of vagueness. Nor can we claim that an adequate theory of meaning is one that would necessarily use the same number of truth values at the referential and at the inferential level. We close, then, by pausing to forestall some objections. First of all, you might note that the notion of validity we have taken from Smith (2008) is not a usual one for fuzzy treatments of vagueness. This is certainly true; for other ways to proceed, see e.g. Machina (1972), Paoli (2003), Smith (2016). But our goal here has not been to provide an overview of fuzzy theories of vagueness; rather, it has been to explore the relations between a particular well-worked-out view—the one of Smith (2008)—and the nontransitive project advocated in our previous work.

Touching on larger issues, you might worry that more than three values may be needed to account for the semantics of comparatives (taking " a is taller than b " to imply that the degree of truth of " a is tall" is greater than that of " b is tall")

¹⁵ We can also see the original four-valued presentation of the strict-tolerant approach in Cobreros et al. (2012) through this lens; the appropriate parameter is $\langle \{0, x, 1-x, 1\}, \{1\}, \{0\} \rangle$, for any $0 < x < 0.5$. These parameters too are proper, symmetric, and open, and so subject to Theorem 10.2.

(Paoli, 2003, 2019), for the semantics of modifiers such as “very” or “determinately” (Lakoff, 1973), or to model degrees of closeness to clear cases (Edgington, 1997; Decock & Douven, 2014). We haven’t shown that three values are sufficient for such purposes, but that was not the goal; what counts as a minimal number of truth values for a fully adequate total semantic theory remains a broader issue.¹⁶

For now, we have at least shown that some of the fuzziness in Smith’s original approach can be shaved with Occam’s razor.

Acknowledgments This paper has a long history and we are grateful to the editors for this opportunity to eventually publish it. The research leading to these results has received support from the Agence Nationale de la Recherche (grants TRILOGMEAN ANR-14-CE30-0010-01 and AMBISENSE ANR-19-CE28-0019-01), from the Department of Cognitive Studies of Ecole normale supérieure (grant FRONTCOG ANR-17-EURE-0017), from the NWO funded ‘Language in Interaction’ project, from the ESSENCE-Marie Curie Initial Training network FP7-PEOPLE-2013-ITN), funded by the European Commission (grant agreement no. 607062) and from PLEXUS (Grant Agreement no 101086295), a Marie Skłodowska-Curie action funded by the EU under the Horizon Europe Research and Innovation Programme. Special thanks to audiences in Cerisy-la-Salle where the first sketch of this paper was presented (Cerisy Context Conference 2011), Opole (ESSLLI 2012 workshop on Trivalent Logics and their Applications), NYU (Vagueness seminar held in 2013), Vienna (2014), Torino (2017), and Cagliari (Trends in Logic, 2022). We thank four anonymous referees for comments on different occasions, including one referee for this volume. We also thank Christian Fermüller, Kathrin Glüer-Pagin, Peter Pagin, Francesco Paoli, Jeff Paris, and Graham Priest, for specific exchanges at different stages of this work.

References

- Barrio, E., Rosenblatt, L., & Tajer, D. (2015). The logics of strict-tolerant logic. *Journal of Philosophical Logic*, 44(5), 551–571.
- Beall, J., & van Fraassen, B. C. (2003). *Possibilities and paradox: An introduction to modal and many-valued logic*. Oxford: Oxford University Press.
- Blasio, C., Marcos, J., & Wansing, H. (2017). An inferentially many-valued two-dimensional notion of entailment. *Bulletin of the Section of Logic*, 46(3/4), 233–262.
- Borel, É. (2014 (1907)). An economic paradox: The sophism of the heap of wheat and statistical truths. *Erkenntnis*, 79(5), 1081–1088. Originally published in French in *La Revue du Mois*.
- Burnett, H. (2014). Penumbral connections in comparative constructions. *Journal of Applied Non-classical Logics*, 24(1–22), 35–60.

¹⁶ Of these, the worry we have encountered most frequently is the one to do with comparatives. The topic is too broad for us to tackle in any detail here, but we will make two quick remarks.

First, we take it that comparative sentences involving gradable adjectives, such as “A is younger than B”, can be treated without the need for additional truth values. Degrees of youth don’t have to be degrees of truth. See in particular Klein (1980) for an influential (three-valued) approach based on delineations—with Burnett (2014) for a recent development—, and Kennedy (2007) for an influential (two-valued) approach based on scalar degrees (distinct from degrees of truth).

Second, even if “true” is *itself* a gradable adjective (see Henderson, 2021; Égré, 2021), and so admits of degrees (which is surely contentious), it would not follow that there is any need to wire such degrees into the semantics of atomic predications in full generality.

- Chemla, E., & Égré, P. (2019). Suszko's problem: Mixed consequence and compositionality. *The Review of Symbolic Logic*, 12(4), 736–767.
- Cobreros, P. (2013). Vagueness: Subvaluationism. *Philosophy Compass*, 8(5), 472–485.
- Cobreros, P., Égré, P., Ripley, D., & van Rooij, R. (2012). Tolerant, classical, strict. *Journal of Philosophical Logic*, 41(2), 347–385.
- Cobreros, P., Égré, P., Ripley, D., & van Rooij, R. (2013). Reaching transparent truth. *Mind*, 122(488), 841–866.
- Cobreros, P., Égré, P., Ripley, D., & van Rooij, R. (2015a). Pragmatic interpretations of vague expressions: Strongest meaning and nonmonotonic consequence. *Journal of Philosophical Logic*, 44(4), 375–393.
- Cobreros, P., Égré, P., Ripley, D., & van Rooij, R. (2015b). Vagueness, truth, and permissive consequence. In T. Achourioti, H. Galinon, K. Fujimoto, & J. M. Fernández (Eds.), *Unifying the philosophy of truth* (pp. 409–430). Springer.
- Decock, L., & Douven, I. (2014). What is graded membership? *Noûs*, 48(4), 653–682.
- Edgington, D. (1997). Vagueness by degrees. In R. Keefe, & P. Smith (Eds.), *Vagueness: A reader* (pp. 294–316). Cambridge, MA: MIT Press.
- Égré, P. (2011a). Perceptual ambiguity and the sorites. In R. Nouwen, R. van Rooij, H.-C. Schmitz, & U. Sauerland (Eds.), *Vagueness and communication* (pp. 64–90). Springer.
- Égré, P. (2011b). Review of Nicholas Smith, *Vagueness and degrees of truth*. *Australasian Journal of Philosophy*, 89(1), 177–180.
- Égré, P. (2017). Vague judgment: a probabilistic account. *Synthese*, 194(10), 3837–3865.
- Égré, P. (2021). Half-truths and the liar. In C. Nicolai, & J. Stern (Eds.), *Modes of truth* (pp. 18–40). Routledge.
- Égré, P., & Barberousse, A. (2014). Borel on the heap. *Erkenntnis*, 79, 1043–1079.
- Fara, D. G. (2000). Shifting sands: An interest-relative theory of vagueness. *Philosophical Topics*, 28(1), 45–81. (Originally published under the name “Delia Graff”).
- Fara, D. G. (2001). Phenomenal continua and the sorites. *Mind*, 110(440), 905–935. (Originally published under the name “Delia Graff”).
- Field, H. (2008). *Saving truth from paradox*. Oxford: Oxford University Press.
- Frankowski, S. (2004). Formalization of a plausible inference. *Bulletin of the Section of Logic*, 33, 41–52.
- French, R., & Ripley, D. (2019). Valuations: Bi, tri, and tetra. *Studia Logica*, 107(6), 1313–1346.
- Gaifman, H. (2010). Vagueness, tolerance and contextual logic. *Synthese*, 174(1), 5–46.
- Gómez-Torrente, M. (2011). The sorites, linguistic preconceptions, and the dual picture of vagueness. In R. Dietz & S. Moruzzi (Eds.), *Cuts and clouds* (pp. 228–253). Oxford University Press.
- Haack, S. (1996). *Deviant logic, fuzzy logic*. Chicago: The University of Chicago Press.
- Henderson, J. (2021). Truth and gradability. *Journal of Philosophical Logic*, 50(4), 755–779.
- Humberstone, L. (1988). Heterogeneous logic. *Erkenntnis*, 29(3), 395–435.
- Kamp, H. (1981). The paradox of the heap. In U. Mönnich (Ed.), *Aspects of philosophical logic* (pp. 225–277). D. Reidel.
- Kennedy, C. (2007). Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, 30(1), 1–45.
- Klein, E. (1980). A semantics for positive and comparative adjectives. *Linguistics and philosophy*, 4(1), 1–45.
- Lakoff, G. (1973). Hedges: A study in meaning criteria and the logic of fuzzy concepts. *Journal of Philosophical Logic*, 2(4), 458–508.
- Lassiter, D. (2011). Vagueness as probabilistic linguistic knowledge. In *International Workshop on Vagueness in Communication* (pp. 127–150). Springer.
- Lassiter, D., & Goodman, N. D. (2017). Adjectival vagueness in a Bayesian model of interpretation. *Synthese*, 194(10), 3801–3836.
- Łukasiewicz, J. (1920). [On three-valued logic]. *Ruch Filozoficzny*, 5, 70–71. English translation in J. Łukasiewicz, *Selected Works*, L. Borkowski ed., 1970.
- Machina, K. F. (1972). Vague predicates. *American Philosophical Quarterly*, 9(3), 225–233.

- Malinowski, G. (1990). Q-consequence operation. *Reports on Mathematical Logic*, 24, 49–54.
- Malinowski, G. (2009a). Beyond three inferential values. *Studia Logica*, 92(2), 203–213.
- Malinowski, G. (2009b). Concerning intuitions on logical many-valuedness. *Bulletin of the Section of Logic*, 38(3/4), 111–121.
- Malinowski, J. (2004). Strawsonian presuppositions and logical entailment. *Logique et Analyse*, 47, 123–138.
- Manor, R. (2006). Solving the heap. *Synthese*, 153(2), 171–186.
- Pagin, P. (2011). Vagueness and domain restriction. In P. Égré & N. Klinedinst (Eds.), *Vagueness and language use* (pp. 283–307). Palgrave.
- Paoli, F. (2003). A really fuzzy approach to the sorites paradox. *Synthese*, 134(3), 363–387.
- Paoli, F. (2019). Degree theories and the sorites paradox. In S. Oms & E. Zardini (Eds.), *The sorites paradox* (pp. 151–167). Cambridge University Press.
- Paris, J. B., Muino, D. P., & Rosefield, M. (2009). Inconsistency as qualified truth: A probability logic approach. *International Journal of Approximate Reasoning*, 50(8), 1151–1163.
- Priest, G. (2008). *An introduction to non-classical logic: From if to is* (2nd ed.). Cambridge: Cambridge University Press.
- Priest, G. (2013). Vague inclosures. In Tanaka et al. (2013) (pp. 367–378).
- Raffman, D. (1996). Vagueness and context-relativity. *Philosophical Studies*, 81, 175–192.
- Read, S., & Wright, C. (1985). Hairier than Putnam thought. *Analysis*, 45(1), 56–58.
- Ripley, D. (2013a). Paradoxes and failures of cut. *Australasian Journal of Philosophy*, 91(1), 139–164.
- Ripley, D. (2013b). Sorting out the sorites. In Tanaka et al. (2013) (pp. 329–348).
- Rotstein, C., & Winter, Y. (2004). Total adjectives vs. partial adjectives: Scale structure and higher-order modifiers. *Natural Language Semantics*, 12(3), 259–288.
- Rumfitt, I. (2015). *The boundary stones of thought*. Oxford: Oxford University Press.
- Scharp, K. (2013). *Replacing truth*. Oxford: Oxford University Press.
- Shramko, Y., & Wansing, H. (2011). *Truth and falsehood: An inquiry into generalized logical values* (Vol. 36). *Trends in logic*. Springer.
- Smith, N. J. J. (2005). Vagueness as closeness. *Australasian Journal of Philosophy*, 83(2), 157–183.
- Smith, N. J. J. (2008). *Vagueness and degrees of truth*. Oxford: Oxford University Press.
- Smith, N. J. J. (2016). Consonance and dissonance in solutions to the sorites. In O. Bueno & A. Abasnezhad (Eds.), *On the sorites paradox*. Springer. Forthcoming.
- Sorensen, R. (2001). *Vagueness and contradiction*. Oxford: Oxford University Press.
- Suszko, R. (1977). The Fregean axiom and polish mathematical logic in the 1920s. *Studia Logica*, 36(4), 377–380.
- Takeuti, G. (1987). *Proof theory* (2nd ed.). *Studies in logic 81*. Amsterdam: Elsevier Science.
- Tanaka, K., Berto, F., Mares, E., & Paoli, F. (Eds.). (2013). *Paraconsistency: Logic and applications*. Dordrecht: Springer.
- Tsuji, M. (1998). Many-valued logics and Suszko's Thesis revisited. *Studia Logica*, 60(2), 299–309.
- van Rooij, R. (2011). Implicit versus explicit comparatives. In P. Égré & N. Klinedinst (Eds.), *Vagueness and language use* (pp. 51–72). Palgrave Macmillan.
- Vickers, S. (1989). *Topology via logic*. Cambridge: Cambridge University Press.
- Weber, Z., Ripley, D., Priest, G., Hyde, D., & Colyvan, M. (2014). Tolerating gluts: A reply to beall. *Mind*, 123(491), 813–828.
- Williamson, T. (1994). *Vagueness*. London: Routledge.
- Wright, C. (1975). On the coherence of vague predicates. *Synthese*, 30(3), 325–365.
- Wright, C. (2001). On being in a quandary: Relativism, vagueness, logical revisionism. *Mind*, 110(437), 45–98.
- Wright, C. (2003). Vagueness: A fifth column approach. In J. Beall (Ed.), *Liars and heaps: New essays on paradox* (pp. 84–105). Oxford: Oxford University Press.
- Zardini, E. (2008). A model of tolerance. *Studia Logica*, 90(3), 337–368.