Comparing Some Substructural Strategies Dealing with Vagueness

Pablo Cobreros[1], Paul Egré[2], David Ripley[3], and Robert van Rooij[4](✉)

[1] Department of Philosophy, University of Navarra, Pamplona, Spain
[2] CNRS/ENS, PSL/SCAS, Institut Jean-Nicod, Paris, France
[3] Department of Philosophy, University of Connecticut, Storrs, USA
[4] Institute for Logic, Language and Communication, Amsterdam, The Netherlands
r.a.m.vanrooij@uva.nl

1 Introduction

It is well-known that in combination with further premises that look less controversial, the tolerance principle – the constraint that if Pa holds, and a and b are similar in P-relevant respects, Pb holds as well – leads to contradiction, namely to the sorites paradox. According to many influential views of the sorites paradox (e.g. Williamson 1994), we therefore ought to reject the principle of tolerance as unsound.

There are reasons to think of such a view as too drastic and as missing out on the role that such a principle plays in categorization and in ordinary judgmental and inferential practice. Taking a different perspective, the tolerance principle ought not to be discarded that fast, even when viewed normatively. Instead, it corresponds to what might be called a *soft constraint*, or a *default*, namely a rule that we can use legitimately in reasoning, but that must be used with care.

One family of approaches represents the tolerance principle by a certain conditional sentence, of the form: $Pa \wedge a \sim_P b \rightarrow Pb$, and bestows special properties to the conditional to turn it into a *soft constraint*. One natural option is to use fuzzy logic, where $v_{\mathcal{M}}(A)$ can be anywhere in $[0, 1]$ and $v_{\mathcal{M}}(A \rightarrow B) = Min\{1, 1 - v_{\mathcal{M}}(A) + v_{\mathcal{M}}(B)\}$. One can demand that the tolerance conditional may never have a value below $1 - \varepsilon$ for some small ε. Given an appropriate sorites sequence, it will be possible to have: $v_{\mathcal{M}}(Pa_1 \rightarrow Pa_2) = 1-\varepsilon$, $v_{\mathcal{M}}(Pa_2 \rightarrow Pa_3) = 1 - \varepsilon$, without having $v_{\mathcal{M}}(Pa_1 \rightarrow Pa_3) = 1 - \varepsilon$.

A different option is to treat the tolerance conditional as expressing a *defeasible rule* (like when '$\rightarrow$' expresses a counterfactual conditional). Say that $Pa \wedge a \sim_P b \rightarrow Pb$ is true provided Pb is true in all 'optimal' $(Pa \wedge a \sim_P b)$-worlds. Call a world $(Pa \wedge a \sim_P b)$-optimal if a is P-similar to b but is not close to a borderline case of P. From $Pa \wedge a \sim_P b, Pa \wedge a \sim_P b \rightarrow Pb$, it need not follow that Pb, since a world may satisfy $Pa \wedge a \sim_P b$ without being

Except for the last section, this paper is an abridged version of a longer paper entitled "The Tolerance Principle: Nontransitive Reasoning or Nonmonotonic Reasoning?". We are indebted to two anonymous reviewers for helpful comments.

© Springer International Publishing Switzerland 2016
J.P. Carvalho et al. (Eds.): IPMU 2016, Part II, CCIS 611, pp. 161–172, 2016.
DOI: 10.1007/978-3-319-40581-0_14

$(Pa \wedge a \sim_P b)$-optimal, precisely when b is a borderline case of P. This conditional fails to satisfy modus ponens, and it is also nontransitive. But moreover it is nonmonotonic, since a world that is $Pa \wedge a \sim_P b$-optimal need not be $Pa \wedge Pc \wedge a \sim_P b$-optimal. On that view, the tolerance conditional represents a defeasible rule, usable only if the main premise correspond to an optimal world.

Both the fuzzy approach and the nonmonotonic approach have some appeal. On the definition of fuzzy validity that allows tolerance to be a sound premise, the sorites paradox is solved by saying that modus ponens is not a valid rule any more (where $\Gamma \models A$ iff $\forall \mathcal{M} : Min\{v_{\mathcal{M}}(\gamma) : \gamma \in \Gamma\} \leq v_{\mathcal{M}}(A)$). On the nonmonotonic approach, the sorites paradox is solved by saying that modus ponens is a defeasible rule: the sorites argument has sound premises, but is not undefeasibly valid. Despite that, both approaches suffer an important limitation, which concerns their treatment of the tolerance principle in terms of a special conditional connective. As is well known, a sorites argument can be stated using only conjunction and negation, by saying that it is not the case that there are two cases a and b that are very similar, but are such that Pa and not Pb. But a nonmonotonic treatment of the conditional does not tell us how to address that alternative version of the sorites. Similarly, for the fuzzy case.

In this paper, we are interested in accounts of vagueness that, instead of relying on a special conditional connective in a way that leaves intuitively desirable properties of a conditional connective in place, and in a way suitable to deal with the sorites argument in its conjunctive form as well as its conditional one. We will focus on two structural approaches which mirror the nontransitive and nonmonotone conditional to some extent, but shift those properties up one level, namely to the consequence relation. The first is the *nontransitive* treatment of logical consequence favored in our past work, on which the principle of tolerance comes out as valid in rule form, but cannot be iterated without risk (soft consequence as *permissive consequence*, see Cobreros et al. (2014) for an overview). The second is the *nonmonotone* treatment of logical consequence, on which the principle of tolerance too can come out as valid, but in a way that is sensitive to context and to the addition of further premises (soft consequence as defeasible consequence).

2 Validating Tolerance Using Non-standard Entailment

We have argued in Cobreros et al. (2012, 2014) that the tolerance principle should be adopted both in rule form and in sentential form. We were able to do so by (i) interpreting the language in standard three valued models, (ii) adding similarity relations (one for each predicate P) to the language and interpreting that in a specific way, and (iii) by formulating a new consequence relation, $\models^{st}_{\sim}$. In this section we first rehearse the details of our approach in more detail. We first concentrate on (i) and (iii). We then consider two broader variants of our initial strategy, which both rely on a more general notion of pragmatic entailment which turn out to be non-monotonic.

2.1 The Logic *st*

Let $\mathcal{M} = \langle \mathcal{D}, \mathcal{I} \rangle$, with $\mathcal{I}$ a total function from atomic sentences to $\{0, \frac{1}{2}, 1\}$. This model extends to formulas according to the strong Kleene valuation scheme:

- $\mathcal{V}_{\mathcal{M}}(\phi) = \mathcal{I}_{\mathcal{M}}(\phi)$, if ϕ is atomic
- $\mathcal{V}_{\mathcal{M}}(\neg\phi) = 1 - \mathcal{V}_{\mathcal{M}}(\phi)$
- $\mathcal{V}_{\mathcal{M}}(\phi \wedge \psi) = min\{\mathcal{V}_{\mathcal{M}}(\phi), \mathcal{V}_{\mathcal{M}}(\psi)\}$
- $\mathcal{V}_{\mathcal{M}}(\phi \vee \psi) = max\{\mathcal{V}_{\mathcal{M}}(\phi), \mathcal{V}_{\mathcal{M}}(\psi)\}$
- $\mathcal{V}_{\mathcal{M}}(\forall x\phi) = min\{\mathcal{V}_{\mathcal{M}}([^{x}/_{\underline{d}}]\phi) : d \in D\}$[1]

We say that ϕ is **strictly true** in $\mathcal{M}$ iff $\mathcal{V}_{\mathcal{M}}(\phi) = 1$, and that ϕ is **tolerantly true** iff $\mathcal{V}_{\mathcal{M}}(\phi) \geq \frac{1}{2}$. In terms of this semantics we can define some well-known logics: Kleene's $K3$ and Priest's LP. Both logics understand entailment as preservation of truth in all models, the difference is that while for $K3$ *truth* means *strict truth*, for LP means *tolerant truth*:

$$\Gamma \models_{K3} \Delta \text{ just in case for all } \mathcal{M}:$$
$$\text{if } \forall A \in \Gamma : \mathcal{V}_{\mathcal{M}}(A) = 1, \text{ then } \exists B \in \Delta : \mathcal{V}_{\mathcal{M}}(B) = 1$$
$$\Gamma \models_{LP} \Delta \text{ just in case for all } \mathcal{M}:$$
$$\text{if } \forall A \in \Gamma : \mathcal{V}_{\mathcal{M}}(A) > 0, \text{ then } \exists B \in \Delta : \mathcal{V}_{\mathcal{M}}(B) > 0$$

A fundamental idea in Cobreros et al. (2012) was to define entailment from *s*trict to *t*olerant:

$$\Gamma \models^{st} \Delta \text{ just in case for all } \mathcal{M}:$$
$$\text{if } \forall A \in \Gamma : \mathcal{V}_{\mathcal{M}}(A) = 1, \text{ then } \exists B \in \Delta : \mathcal{V}_{\mathcal{M}}(B) > 0$$

Thus, although we don't give up the idea that entailment is truth-preserving, we allow the standard of assertion of the conclusions to be weaker than the standard of assertion of the premises. A surprising feature of this logic is that although the semantics makes use of three truth-values, the consequence relation is exactly the familiar consequence relation of classical logic for the standard language fragment. This in contrast with $K3$ and LP, which give up many classically valid arguments.

Now, despite its classicality, this new semantics makes room for tolerance without falling prey to the sorites paradox. In order to account for tolerance, we extend the language with similarity relations, $\sim_P$, one for each predicate P.[2] One interpretation is the following:

- $\mathcal{V}_{\mathcal{M}}(a \sim_P b) = 1$ iff $|\mathcal{V}_{\mathcal{M}}(Pa) - \mathcal{V}_{\mathcal{M}}(Pb)| < 1$, 0 otherwise

[1] We assume here for convenience that each $d \in D$ has a name $\underline{d}$.

[2] See Halpern (2008) and van Rooij (2010) for more on the link between vagueness and nontransitive similarity.

The resulting logic $st^{\sim}$ is a conservative extension of classical logic, in the sense that any classically valid argument in the old vocabulary remains valid. In addition, the tolerance formula $(\forall x, y((Px \wedge x \sim_P y) \rightarrow Py))$ becomes valid, as does the tolerance argument: $Pa, a \sim_P b \models^{st}_{\sim} Pb$. The endorsement of tolerance does not lead to paradox, however, since tolerance in $st^{\sim}$ leads to non-transitivity: $Pa, a \sim_P b \models^{st}_{\sim} Pb$ and $Pb, b \sim_P c \models^{st}_{\sim} Pc$ BUT $Pa, a \sim_P b, b \sim_P c \not\models^{st}_{\sim} Pc$.

We felt, and still feel, that this is a very intuitive and appealing treatment of the sorites paradox. The treatment however, comes with the limitation that we should make a distinction about the diagnosis of the sorites paradox depending on its formulation. If we look at the sorites as a step-by-step argument based on the validity of tolerance inferences then, though each tolerance inference is valid, validity breaks when we try to chain these inferences. If we consider the sorites argument with the tolerance formula $(\forall x, y((Px \wedge x \sim_P y) \rightarrow Py))$ as an explicit premise, then although that formula is valid, the resulting argument is valid but unsound (the tolerance formula is valid but cannot automatically be appealed to as a premise; it is not suppressible). In short, although the tolerance formula is valid, according to the logic $st^{\sim}$ we are not in a position to draw on it as a premise without further ado. However, there are contexts in which we would like to assert the tolerance formula, in much the same way there are contexts in which we would like to assert a contradiction; we would like assert, from time to time, *in a tolerant sense*

A number of recent psycholinguistic studies (e.g., Alxatib and Pelletier 2011; Ripley 2011; Egré et al. 2013) show that naive speakers find a logical contradiction like 'John is tall and John is not tall' acceptable in cases where John is a borderline tall man. This seems to show that we need to take account of tolerant truth, since tolerant truth exhibits this exact behavior. However, just relying on the notion of tolerant truth would mean that the assertion 'John is tall' would be acceptable in the same situation. The same experimental evidence shows, however, that this is not the case: 'John is tall' is taken to be acceptable only if John is *really* tall. In terms of our three-valued models this could be modeled by saying that the assertion 'John is tall' is acceptable only if John is strictly tall. Similarly, Serchuk et al. (2011) found that classical tautologies like '$Tj \vee \neg Tj$' are not automatically accepted if John is borderline tall. So making use of tolerant *and* of strict truth (which exhibits this latter behavior) seems required.

The conclusion we draw from the previous discussion (cf. Cobreros et al. 2012) is that we should interpret a sentence strictly if possible, and tolerantly otherwise. This interpretation strategy is in line with Grice's (1967) strategy to account for scalar implicatures. Unfortunately, this interpretation strategy taken at face value gives rise to trouble for more complex sentences. Alxatib et al. (2013) show that we wrongly predict that a sentence like 'Adam is tall and not tall, *or* John is rich' not only entails, but even means that John is strictly rich, although it should not entail this and intuitively should mean that *either* Adam is borderline tall *or* John is strictly rich. In Cobreros et al. (2015)

we responded by providing a more sophisticated pragmatic interpretation rule to strengthen the meaning of a sentence.

2.2 Pragmatic Interpretation

To account for this pragmatic strengthening we make use of truth-makers. We propose that the pragmatic interpretation of ϕ makes one exact truth-maker of ϕ as true as possible. To determine what the truth-makers of a sentence are, we follow van Fraassen (1969). We start with a set of basic *state of affairs*, SOA. It is assumed that for every element $\mathbf{p}$ of SOA there is also its complement $\overline{\mathbf{p}} \in SOA$ for which it holds that $\overline{\overline{\mathbf{p}}} = \mathbf{p}$. For simplicity we assume a close correspondence between atomic sentences of the language and the SOAs: with each literal (atomic sentence or its negation) of the language there corresponds exactly one SOA: the state of affairs that makes this literal true. The set of *facts*, $\mathcal{F}$, is just $\wp(SOA) - \{\emptyset\}$, so any non-empty subset of SOA is thought of as a fact. If $\mathbf{p}, \mathbf{q} \in SOA$, then $\{\mathbf{p}\}$ and $\{\mathbf{q}\}$ are atomic facts, and $\{\mathbf{p}, \mathbf{q}\}$ is a conjunctive fact. A fact is what makes a sentence true. But, of course, a sentence might have more than one truth-maker. Atomic sentence p is not only made true by atomic fact $\{\mathbf{p}\}$, but also by conjunctive fact $\{\mathbf{p}, \mathbf{q}\}$. The former one is a *more minimal* truth-maker than the latter. More interestingly, disjunctive sentences might have several minimal truth-makers. The disjunction $p \vee q$, for instance, has two minimal truth-makers: $\{\mathbf{p}\}$ and $\{\mathbf{q}\}$. What we are after, however, is the notion of the *exact* truth-makers for ϕ. We say that the disjunction $p \vee (p \wedge q)$ – although it has only $\{\mathbf{p}\}$ as its minimal truth-maker – has *two* exact truth-makers: $\{\mathbf{p}\}$ and $\{\mathbf{p}, \mathbf{q}\}$. We can give the following simultaneous recursive definition of the set of exact truth- and falsity-makers of ϕ, $T(\phi)$ and $F(\phi)$, respectively:

$$
\begin{array}{ll}
T(p) = \{\{\mathbf{p}\}\} & F(p) = \{\{\overline{\mathbf{p}}\}\} \text{ for atomic } p. \\
T(\neg\phi) = F(\phi) & F(\neg\phi) = T(\phi). \\
T(\phi \wedge \psi) = T(\phi) \otimes T(\psi) & F(\phi \wedge \psi) \ = F(\phi) \cup F(\psi). \\
= \{X \cup Y | X \in T(\phi), Y \in T(\psi)\} & \\
T(\phi \vee \psi) \ = T(\phi) \cup T(\phi) & F(\phi \vee \psi) \ = F(\phi) \otimes F(\psi).
\end{array}
$$

Notice that according to these rules, $T(p) = \{\{\mathbf{p}\}\}$, $T(\neg p) = \{\{\overline{\mathbf{p}}\}\}$, $T(p \vee q) = \{\{\mathbf{p}\}, \{\mathbf{q}\}\}$ and $T(p \wedge q) = \{\{\mathbf{p}, \mathbf{q}\}\}$. We analyse conditionals like $\phi \rightarrow \psi$ as material implication, that is $p \rightarrow q \equiv \neg p \vee q$, and thus $T(p \rightarrow q) = \{\{\overline{\mathbf{p}}\}, \{\mathbf{q}\}\}$.

To account for quantifiers, we assume that for each n-place predicate P the model contains facts like $\mathbf{Pd_1}, , \cdots \mathbf{d_n}$, with each $\mathbf{d_i} \in D$ an individual. We assume for simplicity that each $\mathbf{d} \in D$ has a unique name $\underline{d}$ in the language.

$$
\begin{array}{llllll}
- & T(P\underline{d_1}, \cdots, \underline{d_n}) & = & \{\{\mathbf{Pd_1}, \cdots, \mathbf{d_n}\}\} & F(P\underline{d_1}, \cdots, \underline{d_n}) & = \{\{\overline{\mathbf{Pd_1}, \cdots, \mathbf{d_n}}\}\} \\
- & T(\forall x\phi) & = & \bigotimes_{d \in D} T(\phi[^x/_{\underline{d}}]) & F(\forall x\phi) & = \bigcup_{d \in D} F(\phi[^x/_{\underline{d}}]). \\
- & T(\exists x\phi) & = & \bigcup_{d \in D} T(\phi[^x/_{\underline{d}}]) & F(\exists x\phi) & = \bigotimes_{d \in D} F(\phi[^x/_{\underline{d}}]).
\end{array}
$$

Observe that $T(\forall x Px) = T(P\underline{a}) \otimes T(P\underline{b}) = \{\{\mathbf{Pa}, \mathbf{Pb}\}\}$, if $D = \{a, b\}$. Similarly, $T(\exists x Px) = T(P\underline{a}) \cup T(P\underline{b}) = \{\{\mathbf{Pa}\}, \{\mathbf{Pb}\}\}$. Notice that facts might

not only be incomplete (neither verify nor falsify a sentence), they might also be inconsistent and both verify and falsify a sentence. Indeed, we have not ruled out facts like $\{\mathbf{Tj}, \overline{\mathbf{Tj}}\}$. Such inconsistent facts are crucial for us to model the meaning of vague sentences, expressing in this case that John is borderline tall.

$T(\phi)$ can be thought of as a fine-grained *semantic* interpretation of ϕ. It can be used to determine its standard truth-conditional meaning as given by possible worlds semantics, if a world is taken to be a maximally consistent conjunctive fact. In that case the standard truth-conditional meaning of ϕ, $[\![\phi]\!]$, can be recovered as the set of worlds in which ϕ has a truth-maker:

- $[\![\phi]\!] \stackrel{def}{=} \{w \in W \mid \exists f \in T(\phi) : f \subseteq w\}$.

For our purposes, we retain the insistence that worlds be maximal, that $\mathbf{p} \in w$ or $\overline{\mathbf{p}} \in w$ for each atomic fact $\mathbf{p}$ and world w. But we allow for worlds to be inconsistent, for some worlds to contain both $\mathbf{p}$ and $\overline{\mathbf{p}}$, for some atomic facts $\mathbf{p}$. This allows us to capture the difference between strict and tolerant satisfaction at a world, connecting this truth-maker semantics to our three-valued *st*-models. For each atomic sentence p and *st*-model, or world, w we define $\mathcal{V}_w(p) = 1$ iff $\mathbf{p} \in w$ and $\overline{\mathbf{p}} \notin w$; $\mathcal{V}_w(p) = 0$ iff $\mathbf{p} \notin w$ and $\overline{\mathbf{p}} \in w$, and $\mathcal{V}_w(p) = \frac{1}{2}$ otherwise.

But we did not introduce truth-makers just to recover notions we already had. Our purpose in introducing truth-makers is to define a notion of pragmatic meaning in terms of which we can strengthen the semantic meaning of a sentence. We have suggested above that although we allow for inconsistencies, we can still pragmatically infer that $\neg p$ is not true from the fact that 'p' is said by a reasoning analogue to those involving scalar implicatures. In linguistic pragmatics it is not uncommon to use *minimal models* (e.g. van Rooij and Schulz 2004) to account for scalar pragmatic implicatures. For us, a minimal model, or world, is one that is minimally inconsistent: it doesn't contain more inconsistencies than required. To model this, we will make use of the following definition, with $v <_f w \quad \text{iff}_{df}$ $\{x \in SOA : x \in f \ \& \ \overline{x} \in v\} \subsetneq \{x \in SOA : x \in f \ \& \ \overline{x} \in w\}$):

- $PRAG(\phi) \stackrel{def}{=} \{w \in W \mid \exists f \in T(\phi) : f \subseteq w \ \& \ \neg\exists v \supseteq f : v <_f w\}$.

PRAG gets many predictions correct: (i) 'John is tall' is pragmatically interpreted to mean that John is strictly tall, (ii) 'John is not tall' is predicted to mean that John is not even tolerantly tall, (iii) 'John is tall and John is not tall' means that John is borderline tall, and (iv) 'John is tall or not tall' means that John is not borderline tall. All these predictions are in accordance with recent experimental results reported by Alxatib and Pelletier (2011), Ripley (2011), Serchuk et al. (2011) and Egre et al. (2013). Furthermore, 'John is tall and not tall, and Mary is rich' is pragmatically interpreted to mean that John is borderline tall and Mary strictly rich, which seems intuitively correct. Finally, 'John is tall and not tall, or Mary is rich' is correctly interpreted as saying that John is borderline tall, or Mary is strictly rich.

Let us go back now to the tolerance principle. How *PRAG* interprets it depends partly on the way we interpret similarity statements. We might use similarity statements to constrain our models at least in the following ways:

(i) $\mathcal{V}_{\mathcal{M}}(a \sim_P b) = 1$ iff $|\mathcal{V}_{\mathcal{M}}(Pa) - \mathcal{V}_{\mathcal{M}}(Pb)| < 1$, 0 otherwise, or
(ii) $\mathcal{V}_{\mathcal{M}}(a \sim_P b) = 1 - |\mathcal{V}_{\mathcal{M}}(Pa) - \mathcal{V}_{\mathcal{M}}(Pb)|$.

Notice that they are incompatible with each other, at least if there are any a, b with $|\mathcal{V}_{\mathcal{M}}(Pa) - \mathcal{V}_{\mathcal{M}}(Pb)| = \frac{1}{2}$.

Option (i) is the one we mentioned in Sect. 2.1, as an example constraint. According to this option that $\mathcal{V}_w(Pa \wedge a \sim_P b) = 1$ does not force it to be the case that $\mathcal{V}_w(Pb) = 1$. In Sect. 2.1 this was used to account for the nontransitivity of the logic *st*. For one of the logics discussed below, however, we will adopt option (ii).

Making use of facts and truth-makers, this means that we should assume that $T(a \sim_P b) = \{\{\mathbf{Pa}, \mathbf{Pb}\}, \{\overline{\mathbf{Pa}}, \overline{\mathbf{Pb}}\}\}$ and $F(a \sim_P b) = \{\{\mathbf{Pa}, \overline{\mathbf{Pb}}\}, \{\overline{\mathbf{Pa}}, \mathbf{Pb}\}\}$. Notice that if $\mathcal{V}_w(Pa) = 1$ and $\mathcal{V}_w(a \sim_P b) = 1$, it follows that $\mathcal{V}_w(Pb) = 1$. To see this using facts, observe that from $\mathcal{V}_w(Pa) = 1$, it follows that $\mathbf{Pa} \in w$ and $\overline{\mathbf{Pa}} \notin w$. Because $a \sim_P b$ is true and not false, it follows that $\mathbf{Pb} \in w$ and $\overline{\mathbf{Pb}} \notin w$, and thus that $\mathcal{V}_w(Pb) = 1$.[3]

2.3 Tolerance and Inference Relations

Our pragmatic interpretation rule can be included in the definition of logical consequence to try to overcome the limitations we pointed out above about the assertability of the tolerance formula in $st^{\sim}$. Consider, in particular, the following notion of pragmatic consequence, $\models^{prt}$, that goes from *pr*agmatically strongest to *t*olerant (see Cobreros et al. 2015):

- $\phi \models^{prt} \psi$ iff $PRAG(\phi) \subseteq [\![\psi]\!]^t$.

Thus, for inference we take into account what is (pragmatically) meant by the premise. The fact that we look at what was meant by the premise means that, even though $\phi \wedge \neg\phi \models^{prt} \phi$, it does not hold that $\phi \wedge \neg\phi \models^{prt} \psi$. Thus, explosion is not valid. In this sense, *prt*-entailment is a type of paraconsistent entailment relation. Notice, moreover, that if we extend the language with our similarity relation, our new consequence relation $\models^{prt}_{\sim}$ validates the tolerance formula just as much as $\models^{st}_{\sim}$ did.

How should we extend $\models^{prt}$ so as to allow for multiple premises? This is somewhat tricky. The first thought that comes to mind is the following:

- $\Gamma \models^{prt} \phi$ iff$_{df}$ $\bigcap_{\gamma \in \Gamma} PRAG(\gamma) \subseteq [\![\phi]\!]^t$

One shortcoming of that definition, however, is that the resulting notion of consequence fails the adjunction property (the property that $\phi, \psi \models \chi$ provided $\phi \wedge \psi \models \chi$). In order to avoid that problem, we introduce the following variant on that definition. We restrict attention to finite sets of premisses, and say that

[3] To be sure, we could make use of truth-makers as well to implement the analysis of similarity relation in (i). To do so, however, is somewhat involved, and we won't go into that here.

a finite set of premisses Γ entails ϕ provided the pragmatic meaning of the conjunction of the premisses entails the tolerant meaning of the conclusion. We note $\bigwedge \Gamma$ the conjunction of the premisses in Γ, and call Prt the corresponding notion of validity, to distinguish it from the former:

- $\Gamma \models^{Prt} \phi \quad \text{iff}_{df} \quad PRAG(\bigwedge \Gamma) \subseteq [\![\phi]\!]^t$

Notice that if we extend the language with our similarity relation, our new consequence relation $\models^{Prt}_{\sim}$ validates the tolerance formula just as much as $\models^{st}_{\sim}$ did. The fact that we look at what was meant by the premisses means that, even though $\phi, \neg\phi \models^{Prt} \phi$, it does not hold that $\phi, \neg\phi \models^{Prt} \psi$. Thus, explosion is not valid. In this sense, Prt-entailment is a type of paraconsistent entailment relation, just like LP.

Distinctive about $\models^{Prt}$ are the following three properties: (i) conjunction elimination is valid, which implies that $p \wedge \neg p \models^{Prt} p$; (ii) $\models^{Prt}_{\sim}$ is nontransitive, if based on similarity relation (i); and (iii) it is nonmonotonic. As for (ii), even if both Pa and $a \sim_P b$ have value 1, it is required that Pb have at least value $\frac{1}{2}$, but not that it have value 1. As for (iii), in contrast with $\models^{st}_{\sim}$, the notion $\models^{Prt}_{\sim}$ is *nonmonotonic* in the sense that if $\phi_1 \models^{Prt}_{\sim} \chi$, it might still be the case that $\phi_1, \phi_2 \not\models^{Prt}_{\sim} \chi$. In particular, $Pa,\ a \sim_P b \models^{Prt}_{\sim} Pb$, but $Pa, a \sim_P b, \neg Pa \not\models^{Prt}_{\sim} Pb$.

There are reasons for which one might be unhappy with $\models^{Prt}$, however. We argued in Sects. 1 and 2 that nontransitivity and nonmonotonicity might be desirable features to account for vagueness. One might wonder, however, whether we need both of these properties. Second, if pragmatic interpretation captures what is *meant* by the speaker, one might wonder whether either conjunct can be inferred from the premisse $Pa \wedge \neg Pa$. With this sentence the speaker wants to impart that a is borderline tall. But if a conjunct like 'Pa' is asserted alone, it is pragmatically interpreted to mean that a is strictly tall, and thus that a is not borderline tall. If we want a consequence relation capturing what can be asserted on the basis of antecedent assertions, the inference from $Pa \wedge \neg Pa$ to Pa should *not* be valid according to such a relation (see Alxatib and Pelletier 2011).

To account for the latter type of consequence relation we therefore define the following inference relation (from *Pr*agmatic to *Pr*agmatic interpretation), again restricting ourselves to finite sets of premisses:

- $\Gamma \models^{PrPr} \phi \quad \text{iff}_{df} \quad PRAG(\bigwedge \Gamma) \subseteq PRAG(\phi)$

Thus, for inference we take into account what is (pragmatically) meant by the premisses and by the conclusion. It follows that $\models^{PrPr}$ does not satisfy conjunction elimination: in particular, $p \wedge \neg p \not\models^{PrPr} p$. Even though $\phi \wedge \psi \not\models^{PrPr} \phi$ for arbitrary ϕ and ψ, still $p \wedge q \models^{PrPr} p$. Important for the analysis of vagueness is that $(p \wedge \neg p) \vee q \not\models^{PrPr} q$. Similarly, $\phi \wedge \neg\phi \not\models^{PrPr} \psi$, that is, explosion is not valid. In this sense, $PrPr$-entailment is again a type of paraconsistent

entailment relation, just like Prt. Likewise, the tolerance inference is valid: $Pa, a \sim_P b \models_{\sim}^{PrPr} Pb$, on either of the two interpretations of the similarity relations discussed in the previous section. Again, $PrPr$ is *nonmonotonic* for even as $\phi_1 \models^{Prt} \chi$, it can happen that $\phi_1, \phi_2 \not\models^{PrPr} \chi$. This is already clear from the fact that even though $p \models^{PrPr} p$, we have that $p, \neg p \not\models^{PrPr} p$. And in context of soritical reasoning, $Pa, a \sim_P b \models_{\sim}^{PrPr} Pb$, but $Pa, a \sim_P b, \neg Pa \not\models_{\sim}^{PrPr} Pb$.

3 Comparisons

With the machinery introduced in Sect. 2 we have come to define three different consequence relations: $\models^{st}, \models^{Prt}$ and $\models^{PrPr}$. In Sect. 2.2 we presented a way to capture similarity relations, in order to be able to express tolerance. In these logics tolerance is *internalized* since, in fact, the inference from 'Pa' and '$\mathsf{a} \sim_{\mathsf{P}} \mathsf{b}$' to '$\mathsf{Pb}$' is valid. In this section we review how these logics deal with sorites arguments.

The logic *st* is based on the idea that premises and conclusions of an argument need not be subject to equal standards of satisfaction. If the premises of an argument are true to some strict standard, it suffices for validity if the conclusion is true to some less strict standard. Intuitively, this will lead to breaches of transitivity and this is precisely what happens, according to this logic, in sorites arguments.

The logic Prt combines two features: pragmatic interpretation for the premises and tolerance for the conclusion. If the premises of an argument are classically satisfiable, the argument is Prt-valid just in case it is *st*-valid. When the premises are not classically satisfiable, pragmatic interpretation enters the scene. The logic is nontransitive, as we would expect by its affinities with *st*. It is also nonmonotonic as pragmatic interpretation of the premises will change the range of models.

The logic $PrPr$, like the previous ones, depends on the existence of different standards of satisfaction. This time however, the driving idea for $PrPr$ is that the validity of an argument should be evaluated in connection with those models that provide the highest standards of satisfaction compatible with the statements contained in the premises and in the conclusion. Intuitively, that the set of models vary with what is in the premises or in the conclusion will lead to breaches of monotonicity and this is precisely what happens, according to this logic, in sorites arguments.

Hence, we have outlined three distinct consequence relations: *st*, Prt, and $PrPr$. Of these, the first is monotonic but nontransitive, the second is both nonmonotonic and nontransitive, and the third is nonmonotonic but transitive. All validate tolerance, in both its theorem and argument forms. So each gives a different approach to capturing tolerance, and each allows the soft status of tolerance to be recognized by failing to obey the full budget of usual structural rules. Which consequence relation, then, do we recommend?

None, or all, depending on how our recommendation will be understood. None of these consequence relations captures what we take to be the full story

in play with vague predicates; each is only one window on the underlying phenomena. In particular, we have offered a nonstandard (but reasonably familiar) model-theoretic semantics, on which the models involve two distinct notions of satisfaction, plus a pragmatic story involving van-Fraassen-style truthmakers and a particular bias towards stronger interpretations. No single consequence relation will capture the full texture of this story, but each can reflect something important about it.

For example, consider the data reported by Alxatib and Pelletier (2011), where some participants accept the claim that a borderline case is both tall and not tall, while rejecting the claim that he is tall and rejecting the claim that he is not tall. That is, there are situations in which these participants accept $Ta \wedge \neg Ta$ but reject both Ta and $\neg Ta$. This suggests that a consequence relation in which $Ta \wedge \neg Ta$ does not entail either Ta or $\neg Ta$ may capture some patterns in these participants' judgments; indeed, the relation $PrPr$ invalidates these arguments.

Importantly, however, we do not hold up $PrPr$ for this use simply because its pattern of entailments tracks (some) speakers' behavior. We have in fact offered a model of the pragmatic processes underlying these speaker judgments, and it is $PrPr$ that is sensitive to the outputs of this model both in its premises and in its conclusions. Since speaker judgments are sensitive to pragmatic processes, we predict that speakers will judge in ways that accord with $PrPr$-validity.

On the other hand, it is hard not to think that the classical tradition is on to something, in taking conjunctions to entail their conjuncts, and in other ways besides. $PrPr$, of course, does not reflect this; there are cases in which conjunctions are correctly assertible without either of their conjuncts being so. But since *st* validates every classically-valid argument, we have a story available about just what it is that classical logic gets right: any classically-valid argument whose premises are strictly satisfied must have some conclusion that is tolerantly satisfied.

Again, though, we do not hold up *st* for this use simply because it agrees with classical logic. Rather, *st* is fully semantic, involving only the notions of strict and tolerant satisfaction; the pragmatic part of our apparatus does not enter into it. Classical logic, of course, is much more plausible as a logic for semantics than as a logic for pragmatics;[4] the unified picture we have given respects that.

The logic Prt, finally, gives a picture about an interesting mixed phenomenon: when the premises of a Prt-valid argument are correctly assertible, then some conclusion must be at least tolerantly satisfied.

We have given the raw materials to define nine different (two-sided) consequence relations. None of these is itself the full story; they all reveal different interactions between strict satisfaction, tolerant satisfaction, and the pragmatic processes we have outlined. Some, like $K3$, fail to validate tolerance in any form; it is part of our theory, then, that tolerance (the formula) is not always strictly satisfied, and that the tolerance argument is not guaranteed to preserve strict satisfaction. Overall, then, our approach not only explains *how* tolerance can be valid without the sorites wreaking disaster, it also gives a detailed picture of *the ways in which* tolerance is valid.

[4] The point goes back at least to Grice (1967).

4 Outlook

To finish this paper, we will suggest that by making use of our rule of pragmatic interpretation $PRAG$, and the consequence relation $\models^{PrPr}$ we can account for a pragmatic solution to the sorites paradox, and solve an expressive limitation problem for the theory of transparent truth.

According to this radical pragmatic solution to the problem (Gaifman 2010; Pagin 2010, 2011; Rayo 2010 and van Rooij 2010, arguably all based on Wittgenstein 1953), we can appropriately use a predicate P in a context if and only if it helps to clearly demarcate the set of objects that have property P from those that do not (the *gap-hypothesis*). This solution seems natural: the division of the set of all relevant objects into those that do have property P and those that do not is (i) easy to make and (ii) worth making. In those circumstances, the tolerance principle $(\forall x, y((Px \wedge x \sim_P y) \rightarrow Py))$ will not give rise to inconsistency, but serves its purpose quite well. Only in exceptional situations — i.e., when we are confronted with long sequences of pairwise indistinguishable objects — do things go wrong.

Unfortunately, there is a major problem with this approach: the gap hypothesis seems too strong: Even if there is no clear demarcation between the bigger and the smaller persons of the domain, certainly the tallest person can be called 'tall'. Thus, the gap-hypothesis doesn't seem to allow for such exceptional situations. Fortunately, it does, if we make use of our nonmonotonic consequence relation $\models^{PrPr}$. Notice first that the following holds: $Px, \neg Py \models^{PrPr} x \not\sim_P y$. This basically says that if we have a sequence going from truth value 1 to 0, you expect this to be due to a gap (a pair x_i, x_j such that $x_i \not\sim_P x_j$). Similarly, it holds that if you explicitly say that $x \sim_P y$ (and do not say much more) then you expect that Px and Py have the same truth value. Still, this expectation can be cancelled if it is explicitly said that another individual in the transitive closure of the similarity relation (of course, the similarity relation is only transitively closed with respect to strict truth) doesn't have property P. This cancellation holds if we use our nonmonotonic consequence relation $\models^{PrPr}$. Thus, $Px_1, \neg Px_n,\ x_1 \sim_P x_2, \cdots, x_{n-1} \sim_P x_n \not\models^{PrPr} Px_n$.

Consider, finally, the extension of the language with a transparent truth predicate (and the possibility of self-reference). One can express within the language the truth-conditions of sentences of that language. Unfortunately, this normally also immediately gives rise to the Liar paradox. In Cobreros et al. (2013) it is shown, however, that this problem can be solved, making use of the non-transitive consequence-relation $\models^{st}$. There, however, we had to put limits on expressibility: paradox reappears if a sentence can say that it is, e.g. *strictly true*.

Is there perhaps a way of communicating that a sentence is *only true* or *only false* without explicitly saying it? Priest (2006) proposes that to communicate that a sentence is 'only true' or 'only false', a speaker either makes use of an independent speech act of denial (or rejection), or relies on a conversational implicature. In this paper we can suggest the latter: the pragmatic reasoning from the assertion that ϕ is true, to the conclusion that ϕ is *only* true

(and not also false) can be seen as an implicature, and immediately follows from the pragmatic interpretation of ϕ as here proposed.

References

Alxatib, S., Pelletier, F.R.: The psychology of vagueness: borderline cases and contradictions. Mind Language. **26**, 287–326 (2011)

Alxatib, S., Pagin, P., Sauerland, U.: Acceptable contradictions: pragmatics or semantics? a reply to Cobreros et al. J. Philos. Logic **42**, 619–634 (2013)

Cobreros, P., Egré, P., Ripley, D., van Rooij, R.: Tolerant, classical, strict. J. Philos. Logic **41**, 347–385 (2012)

Cobreros, P., Egré, P., Ripley, D., van Rooij, R.: Reaching transparent truth. Mind **122**, 841–866 (2013)

Cobreros, P., Egrá, P., Ripley, D., van Rooij, R.: Vagueness, truth and permissive consequence. In: Achourioti, T., Galinon, H., Fernández, J.M., Fujimoto, K. (eds.) Unifying the Philosophy of Truth. Logic, Epistemology, and the Unity of Science, vol. 36, pp. 409–430. Springer, Netherlands (2015)

Cobreros, P., Egré, P., Ripley, D., van Rooij, R.: Pragmatic interpretations of vague expressions. J. Philos. Logic **44**, 375–393 (2015)

van Fraassen, B.: Facts and tautological entailment. J. Philos. **66**, 477–487 (1969)

Gaifman, H.: Vagueness, tolerance and contextual logic. Synthese **174**, 5–46 (2010)

Grice, H.P.: Logic and conversation, William James Lectures, Harvard University, Reprinted in Studies in the Way of Words, 1989. Harvard University Press, Cambridge (1967)

Halpern, J.: Vagueness and intransitivity. Rev. symbolic Logic **1**(4), 530–547 (2008)

Pagin, P.: Vagueness and central gaps. In: Dietz, R., Moruzzi, S. (eds.) Cuts and Clouds, pp. 254–272. Oxford University Press, Oxford (2010)

Pagin, P.: Vagueness and domain restriction. In: Egré, P., Klinedinst, N. (eds.) Vagueness and Language Use, pp. 283–307. Palgrave MacMillan, Basingstoke (2011)

Priest, G.: Minimally inconsistent LP. Studia Logica **50**, 321–331 (1991)

Priest, G.: In: Contradiction, 2nd edn. Oxford University Press, Oxford (2006)

Rayo, A.: A metasemantic account of vagueness. In: Dietz, R., Moruzzi, S. (eds.) Cuts and Clouds: Vagueness, its Nature and its Logic, pp. 23–45. Oxford University Press, Oxford (2010)

Ripley, D.: Contradictions at the borders. In: Nouwen, R., van Rooij, R., Sauerland, U., Schmitz, H.C. (eds.) ViC 2009. LNCS (LNAI), vol. 6517, pp. 169–188. Springer, Heidelberg (2011)

van Rooij, R.: Vagueness and linguistics. In: Ronzitti, G. (ed.) Vagueness: A Guide. Logic, Epistemology, and the Unity of Science, vol. 19, pp. 123–170. Springer, Netherlands (2011)

van Rooij, R., Schulz, K.: Exhaustive interpretation of complex sentences. J. Logic, Lang. Inf. **13**, 491–519 (2004)

Serchuk, P., Hargreaves, I., Zach, R.: Vagueness, logic and use: four experimental studies on vagueness. Mind Lang. **26**, 540–573 (2011)

Wittgenstein, L.: Philosophische Untersuchungen/Philosophical Investigations. Blackwell, Oxford (1953). Translated by Anscombe, G.E.M