

Pragmatic Interpretations of Vague Expressions: Strongest Meaning and Nonmonotonic Consequence

Pablo Cobreros · Paul Egré · Dave Ripley ·
Robert van Rooij

Received: 1 October 2013 / Accepted: 2 July 2014
© Springer Science+Business Media Dordrecht 2014

Abstract Recent experiments have shown that naive speakers find borderline contradictions involving vague predicates acceptable. In Cobreros et al. (*Journal of Philosophical Logic*, 41, 347–385, 2012a) we proposed a *pragmatic* explanation of the acceptability of borderline contradictions, building on a three-valued semantics. In a reply, Alxatib et al. (*Journal of Philosophical Logic*, 42, 619–634, 2013) show, however, that the pragmatic account predicts the wrong interpretations for some examples involving disjunction, and propose as a remedy a *semantic* analysis instead, based on fuzzy logic. In this paper we provide an explicit global pragmatic interpretation rule, based on a somewhat richer semantics, and show that with its help the problem can be overcome in pragmatics after all. Furthermore, we use this pragmatic interpretation rule to define a new (nonmonotonic) consequence-relation and discuss some of its properties.

Keywords Vagueness · Pragmatics · Partial logic · Non-monotonic logic

P. Cobreros
University of Navarra, Pamplona, Spain
e-mail: pcobreros@unav.es

P. Egré
Institute Jean-Nicod (CNRS-EHESS-ENS), Paris, France
e-mail: paul.egre@ens.fr

D. Ripley
University of Connecticut, Storrs, Mansfields, CT, USA
e-mail: davewripley@gmail.com

R. van Rooij (✉)
Institute for Logic, Language and Computation, Universiteit van Amsterdam,
Amsterdam, Netherlands
e-mail: r.a.m.vanrooij@uva.nl

1 Introduction

A number of recent experiments [1, 11, 26, 33] have shown that naive speakers find some logical contradictions acceptable, specifically borderline contradictions involving vague predicates such as *tall*. In [8] (henceforth TCS) we proposed a pragmatic account of the acceptability of borderline contradictions, making use of an independently motivated strongest meaning hypothesis. In a recent reply, [2] (henceforth APS) show, however, that the pragmatic account predicts the wrong interpretations for some examples involving disjunction. They propose as a remedy a semantic analysis instead, based on fuzzy logic, but one where conjunction and disjunction are interpreted as intensional operators.

In this paper we concede that our earlier proposal was inadequate, but argue that new intensional operators are not required. We continue making use of a pragmatic strongest meaning hypothesis, but we introduce an independently motivated and somewhat richer notion of semantic meaning and an explicit general rule for pragmatic interpretation.¹ We argue that by doing so our analysis can still be seen as a pragmatic approach, and show that we can still account for the examples making use of three truth values only. In addition, we propose a pragmatic nonmonotonic consequence relation and show that this consequence relation has some appealing properties, especially for the analysis of vagueness.

This paper is organized as follows. In Section 2 we give a streamlined presentation of the semantic and pragmatic approach advocated in TCS:2012, but simplified for ease of comparison. In this section we also show where the pragmatic proposal predicts wrongly, extending the counterexamples found by APS. In Section 3 we define an explicit pragmatic interpretation rule that can handle the counterexamples, making use of the notion of an exact truth-maker. We discuss a number of examples involving complex sentences and indicate the different predictions made here in comparison to APS. In Section 4 we define a consequence relation based on our new pragmatic interpretation rule, which is not only non-transitive (as was the logical consequence relation in TCS), but nonmonotone as well. We argue that it behaves favorably compared to similar consequence relations such as [24] LP^m .

2 The Old Theory and Where it Goes Wrong

2.1 Strict, Tolerant, and Strongest Meaning

In TCS we said that a sentence can be true in three different ways: strictly, classically, and tolerantly. We suggested that thus a sentence can also *be interpreted* in three different ways. Our pragmatic claim was: interpret as strongly as possible. For the purpose of illustration and easy comparison, we can forget about classical

¹According to the strongest meaning hypothesis, if a sentence can give rise to several closely related meanings, the sentence should be interpreted in the strongest possible way. A similar principle was used by [1] as well.

interpretation. This leaves us with two notions of truth, strict and tolerant, and three truth-values. (See [10] for showing that our account of vagueness can be restated using a three-valued logic.) Let $\mathcal{M} = \langle \mathcal{D}, \mathcal{I} \rangle$. Let \mathcal{I} be a total function from atomic sentences to $\{0, 1, \frac{1}{2}\}$. Now we can define the truth values of sentences following Kleene’s truth tables as follows:²

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

We say that ϕ is **strictly true** in \mathcal{M} iff $\mathcal{V}_{\mathcal{M}}(\phi) = 1$, and that ϕ is **tolerantly true** iff $\mathcal{V}_{\mathcal{M}}(\phi) \geq \frac{1}{2}$, i.e. iff $\mathcal{V}_{\mathcal{M}}(\phi) \neq 0$. A sentence ψ is **st-entailed** by a set of premises $\Gamma, \Gamma \models^{st} \psi$, iff $\forall \mathcal{M} : \text{if } \forall \phi \in \Gamma : \mathcal{V}_{\mathcal{M}}(\phi) = 1, \text{ then } \mathcal{V}_{\mathcal{M}}(\psi) \geq \frac{1}{2}$.

As applied to vagueness, the motivation for the logic is that if Adam is a borderline case of a tall man, the sentence ‘Adam is tall’ will have value $\frac{1}{2}$, meaning that the sentence is tolerantly true, but not strictly so. According to our semantics this means that also the negation of this sentence will now receive value $\frac{1}{2}$, just as the conjunction of these two sentences, ‘ $Ta \wedge \neg Ta$ ’. Of course, if Adam is not a borderline case, either ‘ Ta ’ or ‘ $\neg Ta$ ’ will have value 0, and the conjunction ‘ $Ta \wedge \neg Ta$ ’ will receive value 0 as well.

In two-valued semantics, every sentence has just one interpretation: the set of models in which the sentence is true. Making use of two ways a sentence can be true allows for (at least) two different ways a sentence can be interpreted: the set of models in which the sentence is strictly true, or the set of models in which it is tolerantly true. In practice, however, lexical rules do not tell you to interpret a sentence strictly or tolerantly, but the distinction is made on a pragmatic basis. In TCS we propose that the explanation is that we always interpret a sentence pragmatically in the strongest possible way. This accounts for the experimentally observed acceptability of contradictions at the border, because contradictions like ‘ $Ta \wedge \neg Ta$ ’ can only be interpreted as true when tolerant truth is at stake. In TCS we show that it also accounts for the lower acceptability observed by [33] for classical tautologies of the form ‘ $Ta \vee \neg Ta$ ’ when Adam is borderline tall.

If we abbreviate the set of models where ϕ is strictly and tolerantly true, respectively, by $\llbracket \phi \rrbracket^s$ and $\llbracket \phi \rrbracket^t$, this pragmatic interpretation rule in our case comes down to the following:

- $Prag(\phi) = \llbracket \phi \rrbracket^s$, if $\llbracket \phi \rrbracket^s \neq \emptyset$, $\llbracket \phi \rrbracket^t$ otherwise.

A standard objection to three-valued truth-functional analyses of vagueness has always been (cf. [12]) that it cannot account for so-called penumbral connections: it fails to predict that $Ta \wedge \neg Ta$ and $Ta \vee \neg Ta$ should always be unacceptable

²Notice that the semantics for the connectives coincides with those of [19, 22, 23]. As in TCS, we here give a substitutional semantics for simplicity, but nothing hangs on that.

(because contradictory) and acceptable (because tautological), respectively. But, as noted above, and as discussed in TCS and [9], these sentences are in fact not always unacceptable or always acceptable. Our pragmatic analysis in TCS predicts the experimental observations much better. However, there are other examples of penumbral connections discussed in [12] and [17] where it is claimed that a truth-functional three-valued analysis fails to make the correct predictions. Perhaps the most challenging one—though not mentioned in TCS—is the following. Suppose that it is established in the context that a and b are equally tall, meaning that both Ta and Tb would have the same semantic value in all relevant models. Now suppose that in fact they are both borderline tall individuals, and thus both sentences have the value $\frac{1}{2}$. Now look at the following two conditionals (analyzed as material implications): $Ta \rightarrow Tb$ and $Ta \rightarrow \neg Tb$ (or even $Ta \rightarrow Ta$ and $Ta \rightarrow \neg Ta$; the type of sentences discussed by [34]). Intuitively, the first one is acceptable, but the latter is not. The problem for three-valued analyses—or so it is argued—is that they cannot explain the difference in acceptability. And indeed, truth-functional three-valued analyses cannot account for their difference in acceptability in terms of truth-value, because (at least on the standard Kleene-based truth-tables) both conditionals would receive the same truth-value: $\frac{1}{2}$.

Fortunately, the difference in acceptability of the two sentences can be explained in terms of the strongest meaning hypothesis, together with a very natural pragmatic constraint on the appropriate assertability of indicative conditionals. According to Grice's Maxim of Quality, you can only assert a sentence appropriately if you believe it to be true. We extend this idea by saying that if somebody asserts a sentence we assume by default that the speaker asserted it *strictly* and you can only do so appropriately if you believe it to be strictly true. Another natural pragmatic constraint explicitly defended by [16] deals only with indicative conditionals: it is inappropriate to assert an indicative conditional if you believe the consequent to be true (for it would have been more informative to simply assert the consequent in that case). Strengthening this constraint to our three-valued case we say that it is inappropriate to *strictly* assert an indicative conditional if you believe the consequent to be strictly true. We can represent what one believes by a set of models. Now look again at the two conditionals $Ta \rightarrow Tb$ and $Ta \rightarrow \neg Tb$, i.e., $\neg Ta \vee Tb$ and $\neg Ta \vee \neg Tb$. Observe that although their actual semantic values are the same if a and b are both borderline tall, their strict interpretations are not: while the set of models where $Ta \rightarrow \neg Tb$, i.e., $\neg Ta \vee \neg Tb$, is strictly true contains only models where both a and b are strictly not tall (because it is known that a and b are equally tall), the set of models where $Ta \rightarrow Tb$, i.e., $\neg Ta \vee Tb$, is strictly true contains *in addition* also models where both a and b are strictly tall. But this means that on their respective strongest interpretations, only $Ta \rightarrow Tb$ satisfies the two pragmatic constraints mentioned above; $Ta \rightarrow \neg Tb$ does not because the consequent is already believed to be strictly true if the whole conditional is. Arguably, this is enough to explain why the former, but not the latter, is acceptable, even though both have the same actual semantic value.³

³This account of the conditional problem is new, and not discussed in TCS.

It will be useful for the rest of this paper to show that we can reformulate our above pragmatic interpretation rule making use of a standard strategy for pragmatic interpretation by means of *orderings*. Each sentence has a truth value in each model, with the usual ordering $0 < \frac{1}{2} < 1$, and in terms of this we can define the pragmatic interpretation rule as follows:

- $Prag(\phi) = \{\mathcal{M} \mid \mathcal{V}_{\mathcal{M}}(\phi) > 0 \ \& \ \neg \exists \mathcal{N} : \mathcal{V}_{\mathcal{M}}(\phi) < \mathcal{V}_{\mathcal{N}}(\phi)\}$

If we define the set of models where ϕ is at least tolerantly true as $\llbracket \phi \rrbracket^t$, this definition simplifies to

- $Prag(\phi) = \{\mathcal{M} \in \llbracket \phi \rrbracket^t \mid \neg \exists \mathcal{N} \in \llbracket \phi \rrbracket^t : \mathcal{V}_{\mathcal{M}}(\phi) < \mathcal{V}_{\mathcal{N}}(\phi)\}$.

Equivalently, but perhaps closer in presentation to what we suggested in TCS, we can think of an ordering $<_S$ between models, with S a set of sentences, defined as follows: $\mathcal{M} <_S \mathcal{N}$ iff_{def} $\{\phi \in S : \mathcal{V}_{\mathcal{M}}(\phi) = 1\} \subset \{\phi \in S : \mathcal{V}_{\mathcal{N}}(\phi) = 1\}$. For future use we will define a slightly more general pragmatic interpretation rule $\underline{Prag}(S)$ for a set of sentences $S = \{\phi_1, \dots, \phi_n\}$.

- $\underline{Prag}(S) = \{\mathcal{M} \in \llbracket \phi_1 \wedge \dots \wedge \phi_n \rrbracket^t : \neg \exists \mathcal{N} \in \llbracket \phi_1 \wedge \dots \wedge \phi_n \rrbracket^t : \mathcal{M} <_S \mathcal{N}\}$.

It is easy to see that $\underline{Prag}(\{\phi\}) = Prag(\phi)$, which are both equivalent to the one we suggested in TCS:

- $Prag(\phi) = \llbracket \phi \rrbracket^s$, if $\llbracket \phi \rrbracket^s \neq \emptyset$, $\llbracket \phi \rrbracket^t$ otherwise.

As an aside, it is interesting to notice that our pragmatic interpretation rule is not *ad hoc*, but can be motivated independently. We mentioned already the strongest meaning hypothesis. But once we think of the value $\frac{1}{2}$ as ‘true and false’, the pragmatic reasoning from the assertion that ϕ , to the conclusion that ϕ is true and not false can be seen as an instance of a Gricean implicature. Given that the speaker didn’t say $\neg\phi$ in addition, we pragmatically conclude that $\neg\phi$ is not true, because otherwise the knowledgeable speaker would have said so. Similarly, although an assertion of $\phi \vee \neg\phi$ cannot rule out that both ϕ and $\neg\phi$ are true, we conclude by pragmatic interpretation that this is not the case, because otherwise the speaker would have said so.

2.2 Some Problematic Examples

Let us go back to vagueness and look at some examples.

(1) Adam is tall. *Ta*

Out of context, this sentence can be interpreted strictly (with truth value 1), so it will be interpreted as being strictly true.

(2) Adam is tall or Adam is not tall. *Ta \vee \neg Ta*

By our rule for disjunction, this sentence is predicted to leave only those models where either the one or the other disjunct is strictly true. This prediction is in

accordance with the experimental results of [33], where it is found that a significant proportion of naive speakers (in fact, most) do not accept (2) when Adam is a borderline case of a tall man.

(3) Adam is tall and Adam is not tall. $Ta \wedge \neg Ta$

Alxatib and Pelletier [1] found that naive speakers agreed with sentences like this for borderline cases of ‘tall’ but not for cases that were either clearly tall or clearly not tall [11, 26] found similar results for ‘near’, and color predicates like ‘yellow’, ‘orange’, ‘blue’, and ‘green’, respectively). It is easy to see that on our pragmatic interpretation rule this sentence is not interpreted strictly. It can be interpreted tolerantly, however, and that is indeed predicted: (3) is interpreted such that Adam is borderline tall.

(4) Adam is tall and Adam is not tall, or John is a monkey. $(Ta \wedge \neg Ta) \vee Mj$

This type of example is similar to one of APS to show that TCS predicts wrongly. It seems that TCS predicts that (4) is interpreted as saying that John is a monkey. Notice, however, that if we ignore any model where John is even tolerantly a monkey, we end up with the same interpretation as (3): Adam is a borderline case of a tall man. But how can we ignore the models where John is a monkey? It seems that context does the trick. Context didn’t play any role in the pragmatic interpretation so far, but it is quite straightforward to let it play its part. We can simply limit the models we consider by adding an extra contextual parameter c denoting the common ground represented by a class of models⁴ (where $\llbracket \phi \rrbracket_c^t = \llbracket \phi \rrbracket^t \cap c$):

- $Prag(\phi, c) = \{\mathcal{M} \in \llbracket \phi \rrbracket_c^t \mid \neg \exists \mathcal{N} \in \llbracket \phi \rrbracket_c^t : \mathcal{V}_{\mathcal{M}}(\phi) < \mathcal{V}_{\mathcal{N}}(\phi)\}$,
or with S a set of sentences $\{\phi_1, \dots, \phi_n\}$
- $\underline{Prag}(S, c) = \{\mathcal{M} \in \llbracket \phi_1 \wedge \dots \wedge \phi_n \rrbracket_c^t : \neg \exists \mathcal{N} \in \llbracket \phi_1 \wedge \dots \wedge \phi_n \rrbracket_c^t : \mathcal{M} <_S \mathcal{N}\}$.

(5) Adam is tall and Adam is not tall, or John is rich. $(Ta \wedge \neg Ta) \vee Rj$

In contrast with (4), we assume now that it is not known in context whether John is rich or not. In fact, he can be even strictly rich. This is another type of counterexample provided by APS to our earlier proposal, and indeed, this one is not so easily explained away. Notice that the first disjunct of (5) can at most have truth value $\frac{1}{2}$. Thus, in the models where John is strictly rich, sentence (5) has truth value 1, which is higher than the value it receives in any model where John is not strictly rich, but the first disjunct has its maximal truth value. It follows that (5) is interpreted as stating that John is strictly rich.⁵ Intuitively, however, sentence (5) states that *either* Adam is borderline tall *or* John is strictly rich. Thus, the pragmatic interpretation rule suggested in TCS mispredicts for some complex sentences involving a *disjunction*.

⁴The notion of ‘model’ that we use here is not exactly the standard model-theoretic notion, where interpretation of non-logical constants varies between models. In contrast, the notion of ‘model’ we use in section 4 to define various notions of entailment will be the standard model-theoretic one. We hope this will never give rise to confusion.

⁵Interestingly enough, this is exactly what Priest’s LP^m predicts for this type of example as well (see [5]). We will come back to this later.

Perhaps more obviously—and not explicitly mentioned by APS—the analysis also mispredicts for some complex sentences with *conjunction*. Consider the following example:

$$(6) \quad \text{Adam is tall and Adam is not tall, and John is rich.} \quad (Ta \wedge \neg Ta) \wedge Rj$$

Because of the first conjunction, the third conjunct cannot be forced to be strictly true. Thus, it might be that the whole sentence is only tolerantly true. By the pragmatic rule proposed in TCS it is predicted that the sentence should thus be interpreted tolerantly, including the second conjunct. But this seems to be wrong: we want to interpret the sentence as saying that John is strictly rich. Thus, not only in case of disjunctions, but also with conjunctive sentences does the analysis of TCS predict a too weak reading.

3 Interpreting Borderline Contradictions: Semantics or Pragmatics?

3.1 Truthmakers and Enriched Pragmatic Interpretation

We agree that examples like (5) and (6) are problematic for the pragmatic analysis suggested in TCS. But instead of inferring that in order to account for these examples we need a semantic analysis instead of a pragmatic one, we show that our pragmatic analysis is still feasible, if slightly refined. Although the refinement uses a somewhat richer notion of semantic meaning, we feel that the refinement is minimal enough still to be called a pragmatic analysis.

How should we account for pragmatic interpretation such that we can infer that a sentence is *only true* if we say that a sentence is true and don't say it is false, such that it can solve the above problems with complex sentences involving borderline contradictions? Here is a proposal: *the pragmatic interpretation of ϕ makes at least one exact truth-maker of ϕ as true as possible*. What are the exact truth-makers of ϕ , and how to think of 'as true as possible'?

As for the first question, an exact truth-maker of ϕ will be thought of as a set of literals, and following [14], the set of exact truth-makers of ϕ , $T(\phi)$, can be defined by a simultaneous recursive definition with the set of exact falsity-makers, $F(\phi)$:⁶

- $T(p) = \{\{p\}\} \quad F(p) = \{\{\neg p\}\} \quad \text{for atomic } p.$
- $T(\neg\phi) = F(\phi) \quad F(\neg\phi) = T(\phi).$
- $T(\phi \wedge \psi) = T(\phi) \otimes T(\psi) = \{A \cup B \mid A \in T(\phi), B \in T(\psi)\}.$
 $F(\phi \wedge \psi) = F(\phi) \cup F(\psi).$
- $T(\phi \vee \psi) = T(\phi) \cup T(\psi).$
 $F(\phi \vee \psi) = F(\phi) \otimes F(\psi).$
- $T(\forall x\phi) = \bigotimes_{d \in D} T(\phi[x/d]) \quad T(\exists x\phi) = \bigcup_{d \in D} T(\phi[x/d]).$
 $F(\forall x\phi) = \bigcup_{d \in D} F(\phi[x/d]) \quad F(\exists x\phi) = \bigotimes_{d \in D} F(\phi[x/d]).$

⁶Van Fraassen uses these truth-makers to give a semantics for the notion of 'tautological entailment' introduced by [3], a notion of entailment that is weaker than both Kleene's K3 and Priest's LP. For recent work on truth-makers, see [13]. For use of the same framework for quite a different purpose, see [28].

Notice that according to these rules, $T(p) = \{\{p\}\}$, $T(\neg p) = \{\{\neg p\}\}$, $T(p \vee q) = \{\{p\}, \{q\}\}$, $T(p \vee \neg p) = \{\{p\}, \{\neg p\}\}$, $T(p \wedge \neg p) = \{\{p, \neg p\}\}$, $T((p \wedge \neg p) \vee q) = \{\{p, \neg p\}, \{q\}\}$, $T((p \wedge \neg p) \wedge q) = T(p \wedge (\neg p \wedge q)) = \{\{p, \neg p, q\}\}$, $T(p \wedge (\neg p \vee q)) = \{\{p, \neg p\}, \{p, q\}\}$, and $T((p \vee q) \wedge (r \vee s)) = \{\{p, r\}, \{p, s\}, \{q, r\}, \{q, s\}\}$.⁷

We analyse conditionals like $\phi \rightarrow \psi$ as material implication, and thus as $\neg\phi \vee \psi$. As a result, a sentence of the form $Ta \leftrightarrow \neg Ta$ will have the following set of exact truth-makers: $\{\{Ta, \neg Ta\}\}$. Observe also that $T(p \wedge (p \rightarrow q)) = T(p \wedge (\neg p \vee q)) = T(p) \otimes T(\neg p \vee q) = T(p) \otimes (T(\neg p) \cup T(q)) = \{\{p\}\} \otimes \{\{\neg p\}, \{q\}\} = \{\{p, \neg p\}, \{p, q\}\}$.

It is interesting to notice that $T(\phi)$ can be seen as the semantic value of a sentence. In fact, as a more fine-grained notion than the one we used so-far. From it, we can derive the classical, tolerant and strict semantic meanings. If we think of models as maximally consistent sets of literals, we can define classical meaning simply as follows: $\llbracket \phi \rrbracket^c =_{df} \{\mathcal{M} \mid \exists S \in T(\phi) : S \subseteq \mathcal{M}\}$. We can define tolerant meaning in the same way, though now the models need only be maximal, and need not be consistent. Thus, we have to think of a model \mathcal{M} as a set such that for each atomic sentence p , (i) $p \in \mathcal{M}$, or (ii) $\neg p \in \mathcal{M}$, or (iii) $p, \neg p \in \mathcal{M}$. The strict meaning of a sentence is defined similarly as well, but now the models are consistent sets of literals, although they don't have to be maximal.⁸ But we didn't define $T(\phi)$ simply to determine the notions of meaning we already had. We wanted to use it to determine a pragmatic meaning $PRAG(\phi)$ suitable to account for the strongest interpretation. In order to do so, we have to make sense of the notion 'as true as possible'.

To make sense of this notion, we can make use of the ordering $<_S$ between models mentioned in the previous section, with S a set of literals (now thought of as an exact truth-maker), defined as follows: $\mathcal{M} <_S \mathcal{N}$ iff_{def} $\{\phi \in S : \mathcal{V}_{\mathcal{M}}(\phi) = 1\} \subset \{\phi \in S : \mathcal{V}_{\mathcal{N}}(\phi) = 1\}$.

In terms of these notions we define the pragmatic interpretation of ϕ , $PRAG(\phi, c)$ (where $\llbracket S \rrbracket_c^t$ abbreviates $\llbracket \psi_1 \wedge \dots \wedge \psi_n \rrbracket_c^t$, if $S = \{\psi_1, \dots, \psi_n\}$).

- $PRAG(\phi, c) = \{\mathcal{M} \in \llbracket S \rrbracket_c^t \mid S \in T(\phi) \ \& \ \neg \exists \mathcal{N} \in \llbracket S \rrbracket_c^t : \mathcal{M} <_S \mathcal{N}\}$
 $= \bigcup_{S \in T(\phi)} \underline{Prag}(S, c)$

Notice that for literals and conjunctive sentences, this pragmatic interpretation rule simply tries to make its exact truth-maker as true as possible, i.e., strictly true. Thus, even if one does not *express* that a sentence like 'Adam is tall' is only true, it follows from the above pragmatic interpretation. But in general a sentence might have more than one exact truth-maker, i.e., when the sentence is disjunctive: so the general pragmatic interpretation rule says that it is enough if one of its exact truth-makers is as true as possible. As a result, $PRAG((p \wedge \neg p) \wedge q, c)$ singles out those models in c where p is tolerantly but not strictly true, *and* where q is strictly true, while

⁷Note that the definition of $T(\phi)$ parallels the construction of the disjunctive normal form of ϕ .

⁸To give a semantics of tautological entailments, [14] looks at models that neither have to be maximal, nor consistent.

$PRAG((p \wedge \neg p) \vee q, c)$ singles out those models in c where either p is tolerantly but not strictly true, or where q is strictly true. This is as desired.⁹

3.2 The Semantics of APS

APS's new connectives In contrast to our pragmatic strategy, APS seek to account for acceptable contradictions in a purely *semantic* way. Their approach is based on fuzzy logic, but they introduce two new *modal* connectives. For ease of comparison, we will assume that the semantics is a three-valued one, however, based on the interpretation rules as defined in Section 1.¹⁰ On top of this semantics, they propose two new (modal) connectives: a new conjunction, \wedge , and a new disjunction, \vee . They assume that conjunctive and disjunctive sentences of natural language should always be analysed in terms of these new connectives. The semantics of these connectives is defined in terms of rescaling as follows:¹¹

- $\mathcal{V}_{\mathcal{M}}(\phi \wedge \psi) = \mathcal{V}_{\mathcal{M}}(\phi \wedge \psi)$, if $\mathcal{C}(\phi \wedge \psi) = \mathcal{F}(\phi \wedge \psi)$
 $= \frac{\mathcal{V}_{\mathcal{M}}(\phi \wedge \psi) - \mathcal{F}(\phi \wedge \psi)}{\mathcal{C}(\phi \wedge \psi) - \mathcal{F}(\phi \wedge \psi)}$, otherwise,
 with $\mathcal{C}(\phi) = \max\{\mathcal{V}_{\mathcal{M}}(\phi) | \mathcal{M} \in c\}$ and $\mathcal{F}(\phi) = \min\{\mathcal{V}_{\mathcal{M}}(\phi) | \mathcal{M} \in c\}$
- $\mathcal{V}_{\mathcal{M}}(\phi \vee \psi) = \mathcal{V}_{\mathcal{M}}(\phi \vee \psi)$, if $\mathcal{C}(\phi \vee \psi) = \mathcal{F}(\phi \vee \psi)$
 $= \frac{\mathcal{V}_{\mathcal{M}}(\phi \vee \psi) - \mathcal{F}(\phi \vee \psi)}{\mathcal{C}(\phi \vee \psi) - \mathcal{F}(\phi \vee \psi)}$, otherwise,
 with $\mathcal{C}(\phi)$ and $\mathcal{F}(\phi)$ as defined above.

To see how this works, let's first look at a conjunction with two independent atomic sentences ϕ and ψ , with $\mathcal{V}_{\mathcal{M}}(\phi) = \frac{1}{2}$ and $\mathcal{V}_{\mathcal{M}}(\psi) = 1$. Because $\mathcal{C}(\phi \wedge \psi) = 1 \neq 0 = \mathcal{F}(\phi \wedge \psi)$, the value of $\mathcal{V}_{\mathcal{M}}(\phi \wedge \psi)$ will be $\frac{\min\{\frac{1}{2}, 1\} - 0}{1 - 0} = \frac{1}{2}$, just like in standard multi-valued logic. But now consider the contradiction $Ta \wedge \neg Ta$ at the border, i.e., when $\mathcal{V}_{\mathcal{M}}(Ta) = \frac{1}{2}$. Notice first that $\mathcal{C}(Ta \wedge \neg Ta) = \frac{1}{2} \neq 0 = \mathcal{F}(Ta \wedge \neg Ta)$.

⁹Our pragmatic interpretation rule does not account for the 'scalar implicature' that from ' $(p \wedge \neg p) \vee q$ ' we conclude that only one of the disjuncts is as true as possible. It is easy to change the pragmatic interpretation rule to account for this—and for embedded implicatures—as well (by changing ' $\exists S \in T(\phi)$ ' in the definition $PRAG(\phi, c) = \{\mathcal{M} | \exists S \in T(\phi) : \mathcal{M} \in \llbracket S \rrbracket_c^t \ \& \ \neg \exists N \in \llbracket S \rrbracket_c^t : \mathcal{M} <_S N\}$ which is equivalent to the one used in the main text, into ' $\exists! S \in T(\phi)$ '), but a discussion of this would go beyond the purpose of this paper.

¹⁰Also APS treat conditionals in terms of material implication. Our assumption that we just use 3 truth-values instead of all the ones in $[0, 1]$ has as a consequence that the following sentences $p \wedge \neg p$, $(p \wedge \neg p) \wedge p$ and $(p \wedge \neg p) \wedge \neg(p \wedge \neg p)$ will all have value 1 exactly if p has value $\frac{1}{2}$. In APS, instead, it is predicted that these different sentences have value 1 in different circumstances. For instance, $(p \wedge \neg p) \wedge p$ is predicted to have value 1 iff p has value $\frac{2}{3}$. Notice that on our analysis the three sentences are predicted to be equivalent because they are predicted to have the same exact truth-maker: $\{p, \neg p\}$.

¹¹As it turns out, this rescaling method is virtually identical to the recalibration-method proposed by [18] to account for adjective-noun combinations.

This means that $\mathcal{V}_{\mathcal{M}}(Ta \wedge \neg Ta) = \frac{\min\{\frac{1}{2}, \frac{1}{2}\} - 0}{\frac{1}{2} - 0} = 1$. Thus, contradictions at the border are predicted to have value 1, they are considered to be true. Similarly, it turns out that a classical tautology like $Ta \vee \neg Ta$ receives value 0 in case $\mathcal{V}_{\mathcal{M}}(Ta) = \frac{1}{2}$.

In contrast to what we did in TCS and what we did above, APS do not give a rule how to interpret a sentence. But it is only reasonable to assume that they take the interpretation of a sentence to be the set of all models in which that sentence is true, i.e., receives value 1. Similarly, it seems natural for them to claim that ϕ is *assertable* in a context c if it is true (i.e., receives value 1) in at least one model in c . It follows that if Ta gets value $\frac{1}{2}$ in all models in the context, claiming that a is borderline tall by saying ‘ $Ta \wedge \neg Ta$ ’ is predicted to be unassertable, not because it is redundant to do so, but because it cannot be true.

According to APS, a sentence ψ is entailed by a set of premises Γ iff in all models \mathcal{M} there is at least one premise $\phi \in \Gamma$ such that $\mathcal{V}_{\mathcal{M}}(\phi) \leq \mathcal{V}_{\mathcal{M}}(\psi)$. Notice that because this requires value 1 for theoremhood, this means that the logic has no logical validities, just like Kleene’s K3.¹² They observe that their consequence relation has two remarkable properties: the logic neither validates conjunction elimination nor disjunction introduction. Indeed, because if $\mathcal{V}_{\mathcal{M}}(Ta) = \frac{1}{2}$, it follows that $\mathcal{V}_{\mathcal{M}}(Ta \wedge \neg Ta) = 1$ and $\mathcal{V}_{\mathcal{M}}(Ta \vee \neg Ta) = 0$, from the truth of $Ta \wedge \neg Ta$ it doesn’t follow that any of its conjuncts is true, and $Ta \vee \neg Ta$ doesn’t follow from any of its disjuncts.

Local strengthening There exists a way to think of the proposal of APS as a pragmatic instead of a semantic one after all. On this reinterpretation of their proposal, we don’t introduce two new semantic intensional connectives, but rather make use of *local strengthening*. The difference between our proposal in Section 3.1 and APS on this reinterpretation is not so much one between a pragmatic versus a semantic analysis, but rather one between two pragmatic analyses: one that allows only for *global* pragmatic strengthenings versus one that also allows for *local* strengthening. The global approach is traditionally favoured and (taken to be) motivated by Gricean ideas, but the idea that also (arbitrary) *parts* of a sentence might be strengthened and perhaps even marked by an explicit grammatical strengthening-operator in the logical form of the sentence has recently been defended by several authors, particularly Chierchia, as giving a more adequate theory.¹³

The ‘local pragmatic’ reinterpretation of APS’s proposal is based on the idea that for any formula ϕ and model \mathcal{M} we can define the strengthened interpretation $\mathcal{S}_{\mathcal{M}}(\phi)$ of ϕ in \mathcal{M} as follows:

- $\mathcal{S}_{\mathcal{M}}(\phi) = \mathcal{V}_{\mathcal{M}}(\phi)$, if $\mathcal{C}(\phi) = \mathcal{F}(\phi)$ with $\mathcal{C}(\phi)$ and $\mathcal{F}(\phi)$ as defined above
 $= \frac{\mathcal{V}_{\mathcal{M}}(\phi) - \mathcal{F}(\phi)}{\mathcal{C}(\phi) - \mathcal{F}(\phi)}$, otherwise.

¹²In fact, in case the language does not have special connectives, this consequence relation is S3, the intersection of K3 and LP.

¹³See [6, 7] for empirical arguments. For a rebuttal of some of these claims, see, among others, [15, 27, 30–32].

It turns out that $\mathcal{S}_{\mathcal{M}}(\phi \wedge \neg\phi) = 2x$, if $\mathcal{V}_{\mathcal{M}}(\phi) = x \leq \frac{1}{2}$, and $\mathcal{S}_{\mathcal{M}}(\phi \wedge \neg\phi) = 2(1 - x)$ otherwise. Moreover, $\mathcal{S}_{\mathcal{M}}(\phi \vee \neg\phi) = 1 - \mathcal{S}_{\mathcal{M}}(\phi \wedge \neg\phi)$. Thus in case $\mathcal{V}_{\mathcal{M}}(\phi) = \frac{1}{2}$, $\mathcal{S}_{\mathcal{M}}(\phi \wedge \neg\phi) = 1$ and $\mathcal{S}_{\mathcal{M}}(\phi \vee \neg\phi) = 0$.

Now one could propose that any sentential part ψ of a whole sentence ϕ can be strengthened and that this should be reflected in the logical form of the sentence. The truth-value of ϕ is still functionally dependent on the truth-values of its parts, but if part ψ is marked in the logical form to be strengthened (let's say by $s(\psi)$), the contribution of ψ should be $\mathcal{S}_{\mathcal{M}}(\psi)$ rather than $\mathcal{V}_{\mathcal{M}}(\psi)$. On this proposal, we can account for the problematic examples (5) and (6), by translating them as $s(Ta \wedge \neg Ta) \vee Rj$ and $s(Ta \wedge \neg Ta) \wedge Rj$, respectively. Example (5) is predicted to have value 1 iff either Ta has value $\frac{1}{2}$ or Rj has value 1, and (6) has value 1 iff Ta has value $\frac{1}{2}$ and Rj has value 1.

Although we consider this pragmatic reinterpretation of APS's proposal interesting, we will limit ourselves in the following comparison between our global pragmatic approach and APS to the latter's explicitly endorsed semantic analysis.

3.3 A Comparison

It is easy to see that the examples (1)–(4) receive exactly the same interpretation on both analyses as on our original rules. On the assumption that where we had \wedge and \vee , APS would understand them as their \wedge and \vee , APS would predict the same as well. So let us go immediately to the problematic examples (5) and (6), repeated below.

(5) Adam is tall and Adam is not tall, or John is rich. $(Ta \wedge \neg Ta) \vee Rj$

In contrast with (4), we assume now that it is not known in context whether John is rich or not. It follows that (5) is interpreted as being true iff Adam is a borderline case of a tall man or if John is strictly rich. This is exactly the desired reading according to APS, and the one they also predict (with \wedge and \vee as their \wedge and \vee). Notice that according to our rule the classical equivalence between $(Ta \wedge \neg Ta) \vee Ta$ and Ta is easily seen not to hold: the former is not predicted to mean that Adam is strictly tall, but rather that he is either borderline tall, or strictly tall. The same is predicted by APS.

(6) Adam is tall and Adam in not tall, and John is rich. $(Ta \wedge \neg Ta) \wedge Rj$

In contrast to the analysis of TCS, we now end up with the correct prediction. We predict that Adam is borderline tall and John is strictly rich. The same prediction is made by APS.

Let us now go through the predictions of a few other complex sentences.

(7) It is not the case that (Adam is tall or Adam is not tall). $\neg(Ta \vee \neg Ta)$

This sentence is predicted to be equivalent with $\neg Ta \wedge Ta$, because $T(\neg(Ta \vee \neg Ta)) = F(Ta \vee \neg Ta) = F(Ta) \otimes F(\neg Ta) = F(Ta) \otimes T(Ta) = \{\{-Ta\}\} \otimes \{\{Ta\}\} = \{\{-Ta, Ta\}\}$. The sentence is thus pragmatically interpreted as claiming that Adam is borderline tall, as desired.

(8) It is not the case that (Adam is tall and Adam is not tall). $\neg(Ta \wedge \neg Ta)$

This sentence will only pick out models where either Adam is strictly tall, or Adam is strictly not tall. Again, this is immediate once we realize that $T(\neg(Ta \wedge \neg Ta)) = F(Ta \wedge \neg Ta) = F(Ta) \cup F(\neg Ta) = F(Ta) \cup T(Ta) = \{\{\neg Ta\}\} \cup \{\{Ta\}\}$. APS make the same prediction: (8) can only be true in case Adam is not borderline tall.

For the examples discussed until now, our predictions are the same as the one made by APS. But there are also examples where the predictions are different. Although we don't have strong intuitions about how we interpret these sentences, we still want to mention them. It is hard to have reliable intuitions about such sentences, because they mix English with bracket notation. But we are assuming that appropriate paraphrases in natural language could make the predictions eventually testable.

(9) Adam is tall and not (Adam is tall or John is rich) $Ta \wedge \neg(Ta \vee Rj)$

We predict that the sentence is interpreted as saying that Adam is tolerantly tall and John not even tolerantly rich. APS, on the other hand, predict that the sentence can also be true if Adam is borderline tall, and John borderline rich.

(10) Adam is tall and (Adam is not tall or John is rich) $Ta \wedge (\neg Ta \vee Rj)$

We predict the interpretation where either Adam is borderline tall or where Adam is strictly tall and John strictly rich. APS, on the other hand, predict that Adam is strictly tall and John is strictly rich.¹⁴

4 Pragmatic Entailment

4.1 From Pragmatic Interpretation to Pragmatic Entailment

In TCS we defined the consequence relation \models^{st} as going from strict premises to tolerant conclusions. It was shown that when restricted to the classical vocabulary, our

¹⁴For what it is worth, we can make sense of our predictions for (9) as well as for (10), even though the predictions APS makes for (10) might seem more appropriate. Consider the reading of (10) made salient by the elaboration: 'Adam is (surely) tall, and either he is (also) not tall or else John is rich (but I forget which)'. It seems to us that our prediction for this elaboration is correct, and that APS mispredicts here. Still, even if our prediction turns out to be incorrect, this still wouldn't automatically mean that a semantic analysis would be preferred to a pragmatic one. It is possible to change the pragmatic interpretation rule so as to predict the same interpretation for (10) as APS does. For instance, by assuming that the definition of $T(\phi \wedge \psi)$ should be as given in the main text only if $T(\phi)$ and $T(\psi)$ are singletons. In any other case we should only allow $A \cup B$ to be an element of $T(\phi \wedge \psi)$ for $A \in T(\phi)$ and $B \in T(\psi)$, if $A \cup B$ does not contain a $p!$ that was not yet in either A or B , where $p! = \{p, \neg p\}$:

$$\begin{aligned} T(\phi \wedge \psi) &= \{A \cup B : A \in T(\phi), B \in T(\psi)\}, \text{ if } T(\phi) \text{ and } T(\psi) \text{ are singletons,} \\ &= \{A \cup B : A \in T(\phi), B \in T(\psi), \neg \exists p : p! \subseteq A \cup B \ \& \ p! \not\subseteq A \ \& \ p! \not\subseteq B\}, \text{ otherwise.} \end{aligned}$$

At this point, however, we are undecided on whether we should change the definition of $T(\phi \wedge \psi)$, because we feel that an assertion of a sentence like (10) would be inappropriate because it violates Grice's maxim of Manner. Notice, though, that in case we would interpret not (10), but rather a sentence like $\neg Ta \vee Rj$ in a context where Ta is known to be strictly true, we would already make the same prediction as APS. As it turns out, however, the alternative definition could be relevant for the analysis of pragmatic entailment discussed in the following section.

logic is identical with classical logic, and is more permissive if we add a family of distinguished similarity relations (one for each predicate T) to the language (indicated by going from \models^{st} to \models^{\sim} , where it is assumed that $\mathcal{V}_{\mathcal{M}}(x \sim_T y) \in \{0, 1\}$). Under those assumptions, the tolerance principle $\forall x, y((Tx \wedge x \sim_T y) \rightarrow Ty)$ becomes valid,¹⁵ but the consequence relation, \models^{\sim} , becomes *non-transitive*. The reason is that if we assume that $Tx \wedge x \sim_T y \wedge y \sim_T z$ is strictly true, we can conclude that Ty is at least tolerantly true. And if $Ty \wedge y \sim_T z$ is, or were, strictly true, we could conclude that Tz would be at least tolerantly true. However, the two inferences cannot be joined together: We cannot conclude from the strict truth of $Tx \wedge x \sim_T y \wedge y \sim_T z$ to the tolerant truth of Tz .

Although the non-transitivity of \models^{\sim} in TCS is only limited to cases where the similarity relation plays a role, one might still wonder whether we should not want the consequence relation to be *as transitive as possible* even if the similarity relation plays a crucial role. In whatever way such a consequence relation (call it \models^{spr}) should be accounted for, it would follow that from ' $Tx \wedge x \sim_T y$ ' we would now conclude that Ty is strictly true, and thus could still use it in a subsequent tolerance inference.¹⁶ Notice, though, that such a new consequence relation would be *nonmonotone*, in the sense that if we added ' $y \sim_T z \wedge \neg Tz$ ' as additional premise, we would no longer be allowed to conclude that Ty is strictly true. Of course, we would still be allowed to conclude Ty , but now only tolerantly.

Whether we use ' \models^{\sim} ' or an alternative as suggested above as our consequence relation, explosion (from a contradiction everything can be derived) is predicted to be valid. Although in accordance with classical logic, this is perhaps not as desirable as it may seem if a contradiction like $Ta \wedge \neg Ta$ is interpreted as stating the contingent fact that Adam is borderline tall. An unfortunate consequence of this picture is that the relation between assertion and inference is lost. Some of what was said (or, better, meant) was ignored in determining what can be inferred from it. Using \models^{st} , TCS predict that everything can be derived from $Ta \wedge \neg Ta$, and thus that, in a sense, the sentence is not really treated as an assertion of a contingent proposition that is true (if only tolerantly) in some but not all models.

APS take what is meant into account, and as a consequence relax conjunction elimination. They argue that this is in accordance with the observed data, because from the assertability of ' ϕ and ψ ' it doesn't automatically follow that ' ϕ ' is assertable. Still, a question remains: can't we interpret a sentence like $Ta \wedge \neg Ta$ as saying that Adam is borderline tall, use this information as a premise of an inference, and still preserve conjunction elimination? It turns out that we can, and that it gives rise to an interesting consequence relation. We adopt the following notion of pragmatic consequence \models^{prt} that goes from *pragmatically strongest to tolerant*:

- $\Gamma \models^{prt} \psi$ iff $\bigcap_{\phi \in \Gamma} PRAG(\phi, \mathbf{M}) \subseteq \llbracket \psi \rrbracket_{\mathbf{M}}^t$, with \mathbf{M} the class of all models.

¹⁵because we demand that $x \sim_T y$ is true provided Px and Py have truth values not differing by more than $\frac{1}{2}$.

¹⁶Notice that the simple definition $\Gamma \models^{\sim} \phi$ iff $S(\Gamma) \subseteq PRAG(\phi, S(\Gamma))$ wouldn't work (with $S(\Gamma)$ as the class of models where all the elements of Γ are strictly true) because there will be models where $Tx \wedge x \sim_T y$ is strictly true, but where Ty is only tolerantly true.

Thus, for inference we take into account what is (pragmatically) meant by each premise. According to this notion of entailment, and in contrast to what APS predict, it follows that $\phi \wedge \psi \models^{prt} \phi$ and also $\phi \models^{prt} \phi \vee \psi$, for any ϕ and ψ . The fact that we look at what was meant by the premises 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. On the other hand, this notion coincides with *st*-entailment in case Γ is contradiction-free.

4.2 Main Properties of \models^{prt}

As in most paraconsistent logics, the *Disjunctive Syllogism*, $\phi, \neg\phi \vee \psi \models^{prt} \psi$, (and reading the conditional as material implication, Modus Ponens) does not hold in general, just as in, for instance, [23] LP. The logic LP is very close to the logic based on \models^{st} : it makes use of the same three-valued interpretation function, but defines the consequence relation differently, as preservation of tolerant truth. A sentence ψ is said to be *LP-entailed* by a set of premises Γ , $\Gamma \models^{LP} \psi$, iff $\forall \mathcal{M} : \text{if } \forall \phi \in \Gamma : \mathcal{V}_{\mathcal{M}}(\phi) \geq \frac{1}{2}$, then $\mathcal{V}_{\mathcal{M}}(\psi) \geq \frac{1}{2}$. Both with respect to \models^{LP} and \models^{prt} , modus ponens, or the disjunctive syllogism, fails if ϕ has value $\frac{1}{2}$. Although our new consequence relation does *in general* not allow for the disjunctive syllogistic proof of ψ from $\neg\phi \vee \psi$ and ϕ , it still predicts the inference to go through in case ϕ is non-contradictory and *can* have value 1 (and is, for instance, an atomic sentence of the form ‘John is tall’).¹⁷ In this it differs from LP: in LP one can also have counterexamples when ϕ can but need not have value $\frac{1}{2}$.¹⁸

Another appealing feature of our inference rule is that it validates *conditional proof*: if $\phi \models^{prt} \psi$, then $\models^{prt} \phi \rightarrow \psi$. This is easy to see. Suppose that $\not\models^{prt} \phi \rightarrow \psi$. Then it must be that there is a model in which ϕ has value 1 and ψ value 0. But in that case it doesn’t hold that $\phi \models^{prt} \psi$.¹⁹ The rule of *contraction*, $\phi \rightarrow (\phi \rightarrow \psi) \models \phi \rightarrow \psi$, holds as well. The only way in which $\phi \rightarrow \psi$ is not tolerantly true is when ϕ has value 1, and ψ has value 0. But in that case the premise is false as well. These appealing features \models^{prt} shares with \models^{LP} .

We know from TCS that if we limit ourselves to the classical vocabulary (thus without a similarity relation) \models^{st} is the same consequence relation as that of classical logic, and that the relation \models^{LP} is strictly weaker. It is easy to see that also the

¹⁷Notice that if we had defined pragmatic consequence as follows: $\Gamma \models_c^{prt} \psi$ iff $PRAG(\bigwedge_{\phi \in \Gamma} c) \subseteq \{\mathcal{M} \in c : \mathcal{M} \in \llbracket \psi \rrbracket^t\}$, things would have been different. In that case we could only conclude from $p, \neg p \vee q$ that either p is only tolerantly true, or that q is strictly true. This has an important consequence, though: it can be that $\phi, \psi \models^{prt} \chi$ although $\phi \wedge \psi \not\models^{prt} \chi$, for $p, \neg p \vee q \models^{prt} q$ but $p \wedge (\neg p \vee q) \not\models^{prt} q$. This problem could be solved, however, by adopting the alternative definition of $T(\phi \wedge \psi)$ as given in footnote 14.

¹⁸One might think that an interesting difference with *LP* shows up here: can we not ensure the validity of the disjunctive syllogism $\phi, \neg\phi \vee \psi \models^{prt} \psi$ by adding $\neg(\phi \wedge \neg\phi)$ as an extra premise, even though *LP* cannot (because $\neg(\phi \wedge \neg\phi)$ is an *LP*-tautology)? Unfortunately, adding $\neg(\phi \wedge \neg\phi)$ as an extra premise cannot guarantee the validity of the disjunctive syllogism for \models^{prt} either, as can be seen by taking ϕ to be $p \wedge \neg p$.

¹⁹The other direction doesn’t hold, though. Assume $\models^{prt} \phi \rightarrow \psi$. Now it doesn’t follow that $\phi \models^{prt} \psi$. For take $\phi := p \wedge \neg p$ and $\psi := q$.

relation \models^{prt} is a (proper) subset of \models^{st} . To show that $\Gamma \models^{prt} \phi \Rightarrow \Gamma \models^{st} \phi$, note that any countermodel for \models^{st} is also one where all premises have value 1 and ϕ has value 0. But any such model is also a countermodel for $\Gamma \models^{prt} \phi$. To see that the subset relation is proper, just consider the classical- (and thus also *st*-) valid inference from $p \wedge \neg p$ to q , and the lack thereof according to \models^{prt} . Similarly, we can show that $\Gamma \models^{LP} \phi \Rightarrow \Gamma \models^{prt} \phi$, because any counterexample to the latter is necessarily also a counterexample to the former. The inclusion is proper, because the disjunctive syllogistic proof from $p, \neg p \vee q$ (with atomic p that can have value 1) to q is \models^{prt} -valid, but not \models^{LP} -valid. Thus $\models^{LP} \subset \models^{prt} \subset \models^{st} = \models^{cl}$.

Our new notion of pragmatic consequence has a number of interesting properties. For one thing, this notion *validates* the *tolerance* principle, $\forall x, y[(Px \wedge x \sim_P y) \rightarrow Py]$ —just like the notion \models^{\sim} of TCS—, if we extend the vocabulary with a similarity relation (in that case $\models^{st} \subset \models^{\sim}$ and $\models^{prt} \subset \models^{\sim}$). This immediately follows from the way \models^{prt} is defined, and the fact that with respect to \models^{\sim} , tolerance is validated as well. Equally obvious is that our new consequence relation \models^{\sim} is still *non-transitive*.

4.3 Nonmonotonicity

But in contrast with \models^{st} , we don't have to extend the language to show that \models^{prt} is non-transitive. This gives rise to another interesting use of \models^{prt} . Look at [21] famous 'proof' of explosion, showing that from a contradiction we can derive everything:

- (1) $p \wedge \neg p$
- (2) p from (1) with conjunction elimination
- (3) $p \vee q$ from (2) with disjunction introduction
- (4) $\neg p$ from (1) with conjunction elimination
- (5) q from (3) and (4) with disjunctive syllogism

We have seen that our pragmatic consequence relation \models^{prt} does not allow for explosion, so what, then, is wrong with the above argument? Among those who have disputed the validity of Lewis's 'proof' of explosion, most have questioned either the validity of disjunction introduction or of disjunctive syllogism.²⁰ Obviously, \models^{prt} allows for disjunction introduction. More interestingly, we have seen above that although \models^{prt} does *in general* not allow for the disjunctive syllogistic proof from $\neg\phi \vee \psi$ and ϕ to ψ , it still predicts the inference to go through in case ϕ *can* have value 1 (and is, for instance, an atomic sentence p of the form 'Adam is tall'). But doesn't that mean that in case ϕ is such a contingent sentence that can have value 1, ψ pragmatically follows from $\phi \wedge \neg\phi$? No, it does not, because our notion of pragmatic consequence, \models^{prt} , is not transitive, even if we limit ourselves to classical vocabulary!²¹ It holds that $p \wedge \neg p \models^{prt} p$ and $p \wedge \neg p \models^{prt} \neg p$. It also holds that $p \models^{prt} p \vee q$ (and, in fact, $p \wedge \neg p \models^{prt} p \vee q$). For atomic p one can \models^{prt} -conclude

²⁰Some (e.g. [25]) have also questioned conjunction elimination.

²¹Note though, that $\phi, \neg\phi \models^{prt} \psi$.

from $\neg p$ and $p \vee q$ that q . But one cannot conclude from these \models^{prt} -inferences that $p \wedge \neg p \models^{prt} q$, because the inferences cannot be joined together. The reason is that in case $p \wedge \neg p$ is as true as possible, p can at most be tolerantly true.

One might also say that our pragmatic consequence relation is *context-dependent*, meaning that although the disjunctive syllogism holds in an ‘empty’ context where p might have value 1 — $p, \neg p \vee q \models^{prt} q$ —, this entailment is not preserved in a context where it is known that $p \wedge \neg p$ is true, and thus where p must have value $\frac{1}{2}$, $p \wedge \neg p, \neg p \vee q \not\models^{prt} q$. Indeed, our new consequence relation is *nonmonotonic*: in the sense that it can be that although $\phi_1, \phi_2 \models^{prt} \chi$, it holds that $\phi_1 \wedge \psi, \phi_2 \not\models^{prt} \chi$.²² Another example illustrating that \models^{prt} , or better \models^{\sim} , is nonmonotone is $Ta \wedge a \sim_T b \models^{\sim} Tb$ but $Ta \wedge a \sim_T b \wedge \neg Ta \not\models^{\sim} Tb$.

There are many other paraconsistent logics that are nonmonotonic. One of them is [24] LP^m . LP^m is very similar to the logic based on \models^{prt} , and, as it turns out, it is also non-transitive.²³ Just like our consequence relation, its nonmonotonicity is due to minimizing inconsistency. However, the way inconsistency is minimized is different: in LP^m one looks only at the set of *minimally inconsistent* models where the premises are all tolerantly true. The set of minimally inconsistent models of ϕ as used in LP^m is defined as $MI(\phi) =_{df} \{\mathcal{M} \in \llbracket \phi \rrbracket^t \mid \exists \mathcal{N} \in \llbracket \phi \rrbracket^t : \mathcal{N} < \mathcal{M}\}$, where $\mathcal{N} < \mathcal{M}$ iff $df \{p \in ATOM \mid \mathcal{N} \in \llbracket p \wedge \neg p \rrbracket^t\} \subset \{p \in ATOM \mid \mathcal{M} \in \llbracket p \wedge \neg p \rrbracket^t\}$. Beall [5] notes that LP^m predicts that from a sentence of the form $(\phi \wedge \neg\phi) \vee p$ one can derive p , for the selected models will all be ones where p has value 1. In this sense this is very much like what we predicted in TCS as well. But it was exactly to solve this problem that we constructed our new pragmatic interpretation rule and consequence relation. Thus, we don’t predict that from $(\phi \wedge \neg\phi) \vee p$ one can derive p , showing that our consequence relation \models^{prt} is really different from \models^{LP^m} , $\Gamma \models^{LP^m} \phi \not\Rightarrow \Gamma \models^{prt} \phi$.²⁴

4.4 The Sorites

Another—and for the treatment of vagueness perhaps most appealing—result of using the pragmatic consequence relation is that we don’t need to make a distinction anymore between the Sorites reasoning with and without the tolerance principle as explicit premise. Without the principle as explicit premise it followed in TCS

²²In fact, our consequence relation verifies (slight variants of) all the standard conditions of nonmonotonic reasoning (as stated by [20]): (i) *Reflexivity*: $\phi \models^{prt} \phi$, (ii) *Left Equivalence*: if $PRAG(\phi) = PRAG(\psi)$, then from $\phi \models^{prt} \chi$, it follows that $\psi \models^{prt} \chi$ (Notice that the condition $PRAG(\phi) = PRAG(\psi)$ cannot be replaced by $\models^{prt} \phi \leftrightarrow \psi$.) (iii) *Right Weakening*: if $\llbracket \psi \rrbracket^t \subseteq \llbracket \chi \rrbracket^t$, then $\phi \models^{prt} \chi$ follows from $\phi \models^{prt} \psi$, (Of course not if $\llbracket \psi \rrbracket^t \subseteq \llbracket \chi \rrbracket^t$ were replaced by $\psi \models^{prt} \chi$ or by $\models^{prt} \psi \rightarrow \chi$.) (iv) *Cautious Monotonicity*: if $\phi \models^{prt} \chi$ and $\phi \models^{prt} \psi$, then it follows that $\phi \wedge \psi \models^{prt} \chi$, and (v) *Or*: if $\phi \models^{prt} \chi$ and $\psi \models^{prt} \chi$, then $\phi \vee \psi \models^{prt} \chi$. Obviously, (vi) *Cautious Cut* (if $\phi \wedge \psi \models^{prt} \chi$ and $\phi \models^{prt} \psi$, then $\phi \models^{prt} \chi$) does not hold, at least in case of \models^{\sim} : a \models^{prt} -conclusion indeed may have a lower truth-value than that of the premises it is based on. Notice that in these rules we have always interpreted the premises as one conjunction. For the reason behind that, see footnote 17.

²³Beall [5] notes that $(q \wedge \neg q) \models^{LP^m} (q \wedge \neg q) \vee r$ and $(q \wedge \neg q) \vee r \models^{LP^m} r$, but $(q \wedge \neg q) \not\models^{LP^m} r$.

²⁴In fact, we think that \models^{prt} is strictly weaker than \models^{LP^m} , but leave a discussion of this to another occasion.

that although each step in the argument is valid, the argument as a whole is *invalid*, because the arguments cannot be chained together. We felt, and still feel, that this is intuitively the correct diagnosis of the Sorites paradox. However, in TCS we had to claim that with the tolerance principle as explicit premise, the argument is valid, but that one of the premises (i.e., the tolerance principle) is not true enough to be used as a premise in an \models_{\sim}^{st} sound inference. Making use of the new consequence-relation \models_{\sim}^{prt} , we can also diagnose the Sorites reasoning with the tolerance principle as explicit premise as invalid, even though all the steps are valid. The fact that the tolerance principle $\forall x_i, x_j ((Px_i \wedge x_i \sim_P x_j) \rightarrow Px_j)$ (with $1 \leq i, j \leq n$) cannot be strictly true if both Px_1 and $\neg Px_n$ are taken as premises that are strictly true, does not rule out that it can be used appropriately in a \models_{\sim}^{prt} -inference.

5 Conclusion and Outlook

The main purpose of this paper was to show that by using an independently motivated and somewhat richer notion of semantic meaning, we can change the global pragmatic interpretation rule proposed in TCS so as to handle the counterexamples pointed out by APS to our original account. We believe that this is a significant insight: in contrast to what APS suggest, this shows that we don't need two new *intensional* connectives, or a notion of *local* strengthening, to account for acceptable borderline contradictions.

In this paper we have also introduced a new pragmatic nonmonotonic consequence relation, $\models_{(\sim)}^{prt}$, and observed that it has some interesting properties. As far as the treatment of vagueness is concerned, and in contrast to \models_{\sim}^{st} , the relation \models_{\sim}^{prt} allows us to take the assertion of the tolerance principle and claims of borderline contradictions at face value, and still use them as substantial premises in our reasoning. We think that the use of $\models_{(\sim)}^{prt}$ might be beneficial in other areas as well,²⁵ and that it differs in interesting ways from other paraconsistent nonmonotonic consequence relations such as those discussed in [4]. More interestingly, perhaps, we have not yet made full use of the pragmatic machinery developed in section 3.1 of this paper. Our notion \models^{prt} only looks for the pragmatic interpretation of the premises. What would result if we also interpreted the conclusion as strongly as possible? Two new consequence relations immediately come to mind: if we abbreviate $\bigcap_{\phi \in \Gamma} PRAG(\phi, \mathbf{M})$ by $PRAG(\Gamma, \mathbf{M})$ (with \mathbf{M} the class of all models), these are $\Gamma \models^{pprs} \psi$ iff_{df} $PRAG(\Gamma, \mathbf{M}) \subseteq PRAG(\psi, \mathbf{M})$ and the more dynamic $\Gamma \models^{pprd} \phi$ iff_{df} $PRAG(\Gamma, \mathbf{M}) \subseteq PRAG(\psi, PRAG(\Gamma, \mathbf{M}))$. The main difference between these two consequence relations is that while \models^{pprd} (like \models^{prt}) does satisfy conjunction-elimination, \models^{pprs} does not. In this respect, \models^{pprs} is closer to the consequence relation preferred by APS in that it preserves assertability. It would be interesting to delve deeper into the properties of these consequence relations. We hope to do this, and more, in further work.

²⁵Van Rooij [29] uses this notion in relation to the well-known *knowability*-paradox.

Acknowledgments We wish to thank the reviewers for their helpful comments. Thanks also to audiences and colleagues at the universities of Amsterdam, Tilburg, NYU, and Barcelona where part of this material was presented, for helpful feedback. Financial support for this work was provided by the Marie Curie Initial Training Network ESSENCE-project, the NWO-sponsored ‘Language in Interaction’-project, and the project ‘Logicas no-transitivas. Una nueva aproximacion a las paradojas’, funded by the Ministerio de Economía y Competitividad, Government of Spain. Thanks also to grants ANR-10-LABX-0087 IEC and ANR-10-IDEX-0001-02 PSL, as well as to the the European Research Council (FP7/2007-2013) under ERC Advanced Grant agreement number 229 441-CCC (PI: Recanati).

References

1. Alxatib, S., & Pelletier, F.R. (2011). The psychology of vagueness: borderline cases and contradictions. *Mind and Language*, 26, 287–326.
2. Alxatib, S., P. Pagin, U. Sauerland (2013). Acceptable contradictions: pragmatics or semantics? A reply to Cobreros et al. *Journal of Philosophical Logic*, 42, 619–634.
3. Anderson, A.R., & N.D. Belnap (1962). Tautological entailments. *Philosophical Studies*, 13(1–2), 9–24.
4. Batens, D. (2000). A survey of inconsistency-adaptive logics. In D. Batens et al (eds.), *Frontiers of Paraconsistent Logic* (pp. 49–73). Baldock: Research Studies Press.
5. Beall, Jc. (2013). Why Priest’s reassurance is not reassuring. *Analysis*, 72(3), 517–525.
6. Chierchia, G. (2004). Scalar implicatures, polarity phenomena and the syntax/pragmatics interface. In *Structures and Beyond* (pp. 39–103). Oxford: Oxford University Press.
7. Chierchia, G., Fox, D., Spector, B. (2012). The grammatical view of scalar implicatures and the relationship between semantics and pragmatics. In P. Portner, C. Maienborn, K. von Heusinger (Eds.), *Semantics: An international handbook of natural language meaning*. Berlin: Mouton de Gruyter.
8. Cobreros, P., Egré, P., Ripley, D., van Rooij, R. (2012a). Tolerant, classical, strict. *Journal of Philosophical Logic*, 41, 347–385.
9. Cobreros, P., Egré, P., Ripley, D., van Rooij, R. (2012b). Tolerance and mixed consequence in a S’valuationist setting. *Studia logica*, 100, 853–876.
10. Cobreros, P., P. Egré, D. Ripley, R. van Rooij. (2013). Vagueness, truth and permissive consequence. In Achourioti et al. (Eds), *Unifying the philosophy of truth*. Springer.
11. Egré, P., V. de Cardelle, D. Ripley (2013). Vagueness and order effects in color categorization. *Journal of Logic, Language and Information*, 265–300.
12. Fine, K. (1975). Vagueness, truth, and logic. *Synthese*, 30, 265300.
13. Fine, K. (2014). Truth-Maker semantics for intuitionistic logic. *Journal of Philosophical Logic*, 43, 549–577.
14. van Fraassen, B. (1969). Facts and tautological entailments. *Journal of Philosophy*, 66, 477–487.
15. Franke, M. (2011). Quantity implicatures, exhaustive interpretation, and rational conversation. In *Semantics and Pragmatics* (vol. 4, pp. 1–82).
16. Grice, H.P. (1967). Logic and conversation. In *William James Lectures, Harvard University, reprinted in Studies in the Way of Words, 1989*. Cambridge: Harvard University Press.
17. Kamp, H. (1975). Two theories of adjectives. In E. Keenan (ed.), *Formal Semantics of Natural Language* (pp. 123–155). Cambridge: Cambridge University Press.
18. Kamp, H., & Partee, B. (1995). Prototype theory and compositionality. *Cognition*, 57, 129–191.
19. Kleene, S.C. (1952). In *Introduction to Metamathematics*. Cambridge: MIT Press.
20. Kraus, S., Lehman, D., Magidor, M. (1990). Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 4, 167–207.
21. Lewis, C.I. (1918). In *A Survey of Symbolic Logic*. Calif: Berkeley.
22. Łukasiewicz, J. (1920). O logice trójwartościowej. *Ruch Filozoficzny*, 5, 170–171.
23. Priest, G. (1979). The logic of paradox. *Journal of Philosophical Logic*, 8, 219–241.
24. Priest, G. (1991). Minimally inconsistent LP. *Studia Logica*, 50, 321–331.
25. Read, S. (1988). *Relevant Logic*. Basil Blackwell.
26. Ripley, D. (2011). Contradictions at the borders. In R. Nouwen, R. van Rooij, H.-C. Schmitz, U. Sauerland (Eds.), (2012), *Vagueness in communication*. LICs: Springer.

27. van Rooij, R., & K. Schulz (2004). Exhaustive interpretation of complex sentences. *Journal of Logic, Language and Information*, 13, 491–519.
28. van Rooij, R. (2014). Leibnizian intensional semantics for syllogistic reasoning. In R. Ciuni et al (Eds.), *Recent Trends in Philosophical Logic Trends in Logic* (vol. 41, pp. 179–194). Heidelberg: Springer.
29. van Rooij, R. Nonmonotonicity and Knowability: As Knowable as Possible. In C. Baskent, L. Moss, R. Ramanujam (Eds.), *Festschrift in honor of Rohit Parikh*. Verlag: Springer (to appear).
30. Sauerland, U. (2004). Scalar implicatures in complex sentences. *Linguistics and Philosophy*, 27, 363–391.
31. Schulz, K., & R. van Rooij (2006). Pragmatic meaning and non-monotonic reasoning: the case of exhaustive interpretation. *Linguistics and Philosophy*, 29, 205–250.
32. Spector, B. (2007). Scalar implicatures: exhaustivity & Gricean reasoning. In M. Aloni, & P. Dekker (Eds.), *Questions in Dynamic Semantics Current Issues in the Semantics/Pragmatics Interface 17*. Elsevier.
33. Serchuk, P., Hargreaves, I., Zach, R. (2011). Vagueness, logic and use: Four experimental studies on vagueness. *Mind and Language*, 26, 540–573.
34. Williamson, T. (1994). *Vagueness*. London: Routledge.