Acceptable Contradictions: Pragmatics or Semantics?
A Reply to Cobreros et al.

Abstract

Naive speakers find some logical contradictions acceptable, specifically borderline contradictions involving vague predicates such as *Joe is and isn't tall.* In a recent paper, Cobreros et al. (*Journal of Philosophical Logic* 2011) suggest a pragmatic account of the acceptability of borderline contradictions. We show, however, that the pragmatic account predicts the wrong truth conditions for some examples with disjunction. As a remedy, we propose a semantic analysis instead. The analysis is close to a variant of fuzzy logic, but conjunction and disjunction are interpreted as intensional operators.

1 Introduction

The phenomenon of vagueness has many facets. One important aspect is the existence of borderline cases. We focus on one fact that is particularly interesting from an empirical, linguistic perspective: In some cases, sentences that seem to correspond to contradictions in classical logic are actually not judged contradictory when they involve borderline cases. Example of this type illustrated by (1) have been discussed as least since Kamp & Partee (1995). Recently a number of experimental studies have brought the issue to the forefront (Alxatib & Pelletier 2010; Ripley 2011; Sauerland 2011). These studies confirm that many ordinary speakers are in fact disposed to *accept* contradictions of a certain kind, exemplified by the following:

(1) John is tall and not tall.

Following Ripley, we refer to examples like (1) as *borderline contradictions* in the following. Both Alxatib & Pelletier (2010) and Ripley (2011) show that borderline contradictions have the highest acceptability precisely for borderline cases: e.g. (1) is most acceptable if John's height is $5'11''$ – the borderline height for a Western man to be tall. The question then arises for any view on vagueness how to account for speakers' acceptance of borderline contradictions. Should such dispositions be dismissed as merely confused? Should they be given a semantic account by which contradictions come out true? Should they be given a account in linguistic pragmatics by which they are false but highly *assertible*? Or should we find yet some other way?

In this paper we shall review a pragmatic account of the phenomenon recently offered by Pablo Cobreros, Paul Egré, David Ripley and Robert van Rooij (2011). Cobreros et al. operate with three different notions of satisfaction: *strict*, *classical*, and *tolerant* truth (which inter alia allows them to assert an approximation to an ordinary tolerance principle), and use the interplay between these modes of evaluations for a pragmatic account of borderline contradictions. In section 2, we present the details of the proposal, and show that it faces a problem with sentences that have a borderline contradiction as one part, but contain additional material. Cobreros et al.'s pragmatic proposal would evaluate either the entire sentence strictly or non-strictly. But, we show that there are examples where the intuitively available interpretation is one where the sentence, so to speak, is in part evaluated strictly and in part non-strictly. The pragmatic analysis doesn't allow for this, because it operates at the sentence level. In section 3, we propose a semantic analysis that remedies the problem. Instead of building on a choice of evaluation, the semantic analysis assumes a) a multivalued logic, and b) a semantics of conjunction that is not fully truth-conditional, but contains a modal component: if the two conjuncts are necessarily at least a partial contradiction, their conjunction will be fully true. In section 3, we also discuss further relevant examples that test our proposal, but conclude that the judgment on the examples is impossible to ascertain. Section 4 presents our conclusions.

2 The Pragmatic Proposal for Borderline Contradictions

Cobreros et al. (2011), hereafter CERvR, present a framework that uses three notions of truth for sentences involving vague predicates: in addition to the classical notion of truth, a notion of tolerant truth and a notion of strict truth. They then apply this framework to borderline contradictions. Specifically, they combine their different notions of truth with the Strongest Meaning Hypothesis of Dalrymple et al. (1998), which predicts that a natural language sentence would be evaluated by a speaker relative to the most restrictive evaluation function that would allow it to be true. CERvR show that in basic cases their proposal predicts the acceptability of borderline contradictions. In this section, we briefly introduce their framework and show how it applies to basic borderline contradictions. We then go on to demonstrate that the proposal doesn't extend to more complex sentences involving borderline contradictions.

Consider first the three notions of truth in CERvR, which they attribute to unpublished work by van Rooij. They denote the classical notion of truth with $[\![]\!]^c$, while they use $[\![]\!]^t$ for the tolerant notion of truth and $[\![]\!]^s$ for its dual, the strict notion. The three notions of truth stand in an entailment hierarchy: Strict truth entails classical truth, and classical truth entails tolerant truth. In the formal definitions of strict, classical, and tolerant truth, provided below, the binary relation $\sim_P$ is taken to hold between individuals x and y iff x and y are indistinguishable with respect to their membership in predicate P.

- Classical Truth:
  - (ci) $[\![Pa]\!]^c = 1$ iff $[\![a]\!]^c \in I(P)$
  - (cii) $[\![\neg\phi]\!]^c = 1$ iff $[\![\phi]\!]^c = 0$
  - (ciii) $[\![\phi \vee \psi]\!]^c = 1$ iff $[\![\phi]\!]^c = 1$ or $[\![\psi]\!]^c = 1$
  - (civ) $[\![\phi \wedge \psi]\!]^c = 1$ iff $[\![\phi]\!]^c = 1$ and $[\![\psi]\!]^c = 1$
- Tolerant Truth:
  - (ti) $[\![Pa]\!]^t = 1$ iff $\exists x[x \sim_P [\![a]\!]^c \,\&[\![P]\!]^c(x) = 1]$
  - (tii) $[\![\neg\phi]\!]^t = 1$ iff $[\![\phi]\!]^s = 0$
  - (tiii) $[\![\phi \vee \psi]\!]^t = 1$ iff $[\![\phi]\!]^t = 1$ or $[\![\psi]\!]^t = 1$
  - (tiv) $[\![\phi \wedge \psi]\!]^t = 1$ iff $[\![\phi]\!]^t = 1$ and $[\![\psi]\!]^t = 1$
- Strict Truth:
  - (si) $[\![Pa]\!]^s = 1$ iff $\forall x[x \sim_P [\![a]\!]^c \rightarrow [\![P]\!]^c(x) = 1]$
  - (sii) $[\![\neg\phi]\!]^s = 1$ iff $[\![\phi]\!]^t = 0$
  - (siii) $[\![\phi \vee \psi]\!]^s = 1$ iff $[\![\phi]\!]^s = 1$ or $[\![\psi]\!]^s = 1$
  - (siv) $[\![\phi \wedge \psi]\!]^s = 1$ iff $[\![\phi]\!]^s = 1$ and $[\![\psi]\!]^s = 1$

Clauses (ci-civ) are taken from the standard treatments of first-order predicate logic: an atomic formula Pa is classically true if and only if a denotes an element of $I(P)$, the extension of P relative to model M, and (cii-civ) provide the familiar recursive definitions of the semantics. The basic definition of tolerant truth, (ti), is given in terms of the relation $\sim_P$: a predicate P *tolerantly* holds of an individual constant a if and only if there is an individual x which is indistinguishable from a's value with respect to the applicability of P (that is, $x \sim_P [\![a]\!]^c$), such that Px is classically-true. P *strictly* holds of an individual constant a if *every* individual x that is P-indistinguishable from a is such that Px classically holds. The duality that holds between $[\![]\!]^t$ and $[\![]\!]^s$ is evident from (tii) and (sii): for any formula ϕ, $[\![\phi]\!]^t = 1$ if and only if $[\![\neg\phi]\!]^s = 0$, and $[\![\phi]\!]^s = 1$ iff $[\![\neg\phi]\!]^t = 0$.

First consider the predictions of the proposal for a borderline statement such as (2).

(2) A 5′11″-tall man is tall.

In this case, P would be the predicate *tall* and a a constant denoting an individual of height 5′11″. Since a is a borderline case of tallness by assumption, there must be an x with $x \sim_P [\![a]\!]^c$ and $x \in [\![P]\!]$, but also a y with $y \sim_P [\![a]\!]^c$ and $y \notin [\![P]\!]$. In other words, P holds tolerantly, but not strictly, of a when a is a borderline case of P. The classical truth value of Pa, however, could be either true or false. CERvR now assume a pragmatic principle in addition to the three notions of truth: the Strongest Meaning Hypothesis (SMH) of Dalrymple et al. (1998). Applied to the cases at hand, the SMH states that speakers judge a sentence according to the strongest notion of truth for which there exists a possible scenario that makes it true. The SMH predicts that (2) must be evaluated relative to the strict notion of truth, because there is a possible scenario where Pa is strictly true, namely one where the average height of men is low enough to make 5′11″ definitely tall. In this way, the SMH predicts that (2) should be judged false in the actual scenario where a is a borderline case. This prediction is the desired result for (2): experimental data of Alxatib & Pelletier (2010); Ripley (2011) and others confirm that speakers judge borderline cases such as (2) overwhelmingly false (see also Bonini et al. (1999) for similar findings).

CERvR's proposal also accounts for the difference between acceptable borderline contradictions like (3a) and less acceptable contradictory non-borderline conjunctions such as (3b):

(3) a. A 5′11″ tall man is and isn't tall.

b. A 6′ 4″ tall man is and isn't tall.

For the sake of illustration, continue to assume that a man standing 5′11″ qualifies as a borderline case of tallness, while a height of 6′4″ constitutes a clear case of tallness. On CERvR's system, this means that 5′11″ sits near the border between the classical extension of P and its anti-extension. So, given an individual a whose height is 5′11″, and given that a is standing in a sorites series, there is a P-indistinguishable x from a which belongs to the classical extension of P, and there is also a P-indistinguishable x from a which does not belong to the classical extension P. As just discussed, this qualifies a both as tolerantly-tall *and* as tolerantly not-tall, making $[\![Pa \wedge \neg Pa]\!]^t = 1$. By contrast, if b's height of 6′4″ is great enough to make every indistinguishable individual x a classical case of P, then b qualifies as strictly tall. And owing to the duality that holds between strict truth and tolerant truth, $[\![Pb]\!]^s = 1$ guarantees that $[\![\neg Pb]\!]^t = 0$, and therefore $[\![Pb \wedge \neg Pb]\!]^t = 0$. So for any predicate P, the contradiction $Px \wedge \neg Px$ will be false for those xs of whom P holds strictly, even when the contradiction is evaluated under the tolerant notion of truth. The reason is that whenever one of the individual conjuncts that make up the contradiction holds strictly, the other conjunct, the negation of the first, will be tolerantly false.

Of course, it is only tolerant truth that can validate a contradiction in CERvR's system, and as discussed in the previous paragraph, the contradiction holds tolerantly only for borderline cases. So, to account for the acceptability of contradictions in borderline cases, CERvR invoke the SMH, namely, the provision that among a set of available interpretations (available notions of truth in this case), one is to select the logically strongest interpretation that can be true in at least one model. In the case of contradictions, this is the tolerant notion of truth, for on the stronger two notions, the classical and the strict, a contradiction can never have a value other than false. But in other cases the SMH will impose a stronger requirement. Take the expression Pa for example. As is dictated by the SMH, Pa will be evaluated according to the strongest non-trivially-false notion of truth. Since there are models in which Pa is strictly true, the strongest such notion will be the strict notion. But of course, Pa is strictly true only if a is a strict case of P; if a is a borderline case, then we expect Pa to be unacceptable. So the prediction is that while neither Pa nor $\neg Pa$ are acceptable for a borderline case a, the conjunction $Pa \wedge \neg Pa$ is acceptable. This prediction is actually confirmed in the recent experimental work of Alxatib & Pelletier (2010), whose proposal also relies on the SMH, but in combination with a supervaluationary semantic platform instead.

While CERvR's proposal can account for basic cases of borderline contradictions, the account doesn't extend to some slightly more complex examples. Consider (4), where a borderline contradiction occurs within a disjunction.

(4) Joe either is and isn't tall or Joe has red hair.

CERvR predict (4) to be equivalent to just *Joe has red hair*; because it is logically possible for Joe to have red hair, sentence (4) is not a classical contradiction, and must therefore be evaluated strictly, under the SMH. But because the first disjunct is false under strict evaluation, this makes (4) equivalent to *Joe has red hair*. The prediction, then, is that (4) is judged false if Joe is a *blond* $5'11''$-tall man.

But to the extent that *tall and not tall* is interpreted as *of borderline height*, (4) seems to be equivalent to *Joe is either of borderline tall height or he has red hair.* To see that this is the desired interpretation, consider cases where we conclude from (4) that Joe is of intermediate height once we learn that he couldn't have red hair. For example, if we had to identify Joe in a group of people who all have black hair, and we know that (4) holds, then we would conclude that Joe must be one of the people in the group of intermediate height.

This intuition finds support in experimental studies reported in Sauerland et al. (2011), where MTurk participants were shown an illustration of five men varying in height and also in degree of wealth (the latter indicated using dollar signs, ranging from poor ($) to rich ($$$$$)). The results show a significantly higher rate of assent (45.7%) to (5) when the blond men in the picture are of borderline height, but vary in richness, than when they are all poor but vary in tallness (12.8%).

(5) Every blond man is either rich or tall and not tall.

In the former scenario, all the blond men qualify as 'tall and not tall', since they are each of borderline height, and so they all verify the disjunctive predicate, regardless of their degree of wealth. By contrast, the latter scenario shows only poor men, so none of them satisfy the disjunct 'rich'. The disjunctive predicate can only be satisfied if each of the men satisfies the predicate 'tall and not tall', and since the men are shown to vary in height, this condition is not met.

The reason that (4) and (5) are problematic for CERvR has to do with the pragmatic nature of the proposal. The SMH is a principle of linguistic pragmatics and therefore is applied at the sentence level.[1] Though one of the disjuncts in (4) and (5) is a classical contradiction, the full sentences are not. The SMH predicts, then, that only the strict interpretation can apply and therefore the complex disjuncts, the borderline contradictions, are interpreted as contradictory, contrary to what is suggested by the experimental evidence. It therefore seems necessary to evaluate sentences in part relative to strict evaluation and in part relative to the tolerant notion of truth.

In sum, while the predictions of CERvR's system fit the empirical findings regarding the acceptability of simple vague expressions, including contradictions, the predictions seem to depart from what is observed in cases where the contradictions are themselves disjoined with logically independent propositions. We have seen that, under the SMH, contradictions are evaluated on the tolerant notion of truth, since tolerant truth is the strongest mode of evaluation that can assign contradictions non-false interpretations. But once a logically-independent proposition is introduced as a disjunct to the contradiction, the SMH will favor the strict notion of truth, since for any expressions ϕ and ψ, there will be models in which $[\![(\phi \wedge \neg\phi) \vee \psi]\!]^s = 1$, namely those models in which $[\![\psi]\!]^s = 1$. In the next section we offer a semantic analysis that overcomes this problem.

3 The Semantic Proposal

In the previous section, we discussed a problem with CERvR's proposal. The problem stems from the pragmatic perspective: CERvR apply only one evaluation held constant for an entire sentence. In this section, we develop a semantic alternative that is more flexible. We present our proposal in two parts: It is based on a multivalued logic which we introduce first. The new idea of our proposal is to add to this multivalued logic two new semantic operators, $\circledvee$ and $\circledwedge$, that we intend to model the behavior of natural language disjunction and conjunction better than any $\vee$ and $\wedge$ of multivalued logic.

Multi-valued Logic

The language of the logic under consideration is, like that of basic multivalued logic, identical to first-order predicate calculus. We therefore skip defining the (familiar) language, and proceed to review the semantics of basic

[1] The criticism we offer is directed at CERvR's proposal as it is presented in the main body of their paper. In their footnote 19, however, CERvR suggest the possibility of applying the SMH below the sentence level, though they do not pursue this suggestion in detail. Our proposal may be seen as a way of spelling out the details of this suggestion.

fuzzy logic (e.g. Hajek 2009).[2] We skip the definition of the language since the well-formed formulas are exactly those of classical predicate calculus, and for simplicity we restrict our attention to 1-place predicates only. The interpretation of each predicate letter P is designated as a function $I(P) : D \to [0,1]$. Now,

(6) let v be a function from well-formed formulas to the interval [0,1], then

(i) for any model M, predicate letter P, and term t, $v_M(Pt) = I_M(P)(I(t))$ if t is a constant symbol, and $I_M(P)(g(t))$ if t is a variable symbol (where g is the assignment function)

(ii) $v_M(\neg\phi) = 1 - v_M(\phi)$

(iii) $v_M(\phi \vee \psi) = \max(v_M(\phi), v_M(\psi))$

(iv) $v_M(\phi \wedge \psi) = \min(v_M(\phi), v_M(\psi))$

It is fairly obvious that this multi-valued logic will not predict the acceptability of borderline contradictions, since the highest degree of truth they can reach is 0.5. This feature is not specific to the multi-valued logic just defined, but holds rather generally of truth functional multi-valued logic (see Sauerland 2011 for discussion).

Intensional Conjunction and Disjunction

We add to the logic two new binary operators: $\ovee$ and $\owedge$. These, we propose, capture the properties of disjunction and conjunction in natural language. The definitions of the two operators are based on the observation that for formulas expressing a borderline contradiction $\phi \wedge \neg\phi$ the range of possible truth values is not the full range [0,1], but only the subset [0,0.5]. The operators $\ovee$ and $\owedge$ are defined so as to utilize the full range [0,1] whenever possible. This is accomplished by scaling the range of truth values to extend from 0 to 1 as illustrated in the following graphic. The following graph shows the truth value of a borderline contradiction $a \wedge \neg a$ (broken line) and of $a \owedge \neg a$ (continuous line). We assume in the graph that $\mathcal{M} = [0,1]$ and $V_M(a) = M$ for any $M \in [0,1]$.

[2] We chose to review fuzzy logic for presentational purposes only; the semantics we propose does not depend of the specific form of the multi=valued logic.

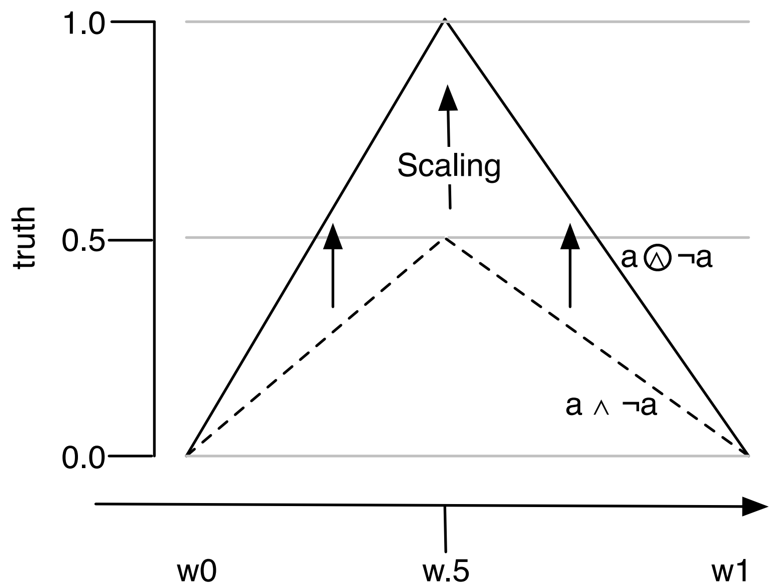

To achieve this scaling effect, we evaluate the acceptability of each formula based not only on its truth-value, but also on how true the formula is *within the range of values that it can have*. In order to do this, we identify for each formula a pair of values that correspond to the truest and the falsest that the formula can be. These we call, respectively, the *ceiling* (**c**) and the *floor* (**f**), and we define them in (C) and (F) below. The definitions use the mathematical concepts of supremum (sup) and infimum (inf) of a set. The supremum of a set S is defined as the smallest number that is greater than or equal to any element of S, and the infimum is the greatest number smaller than or equal to any element of S. Note that if a set S has a maximum and/or a minimum, $\max(S) = \sup(S)$ and $\min(S) = \inf(S)$, but that every bounded set of real numbers has a supremum and an infimum.

(7) (C) $\mathbf{c}(\Phi) = \sup\{k : \text{for some model } M, v_M(\Phi) = k\}$
(F) $\mathbf{f}(\Phi) = \inf\{k : \text{for some model } M, v_M(\Phi) = k\}$

The ceiling (or floor) of a formula Φ is the highest (or lowest) degree of truth it can get: take all the values k for which at least one model M is such that $v_M(\Phi) = k$, then the greatest such k is said to be the ceiling of Φ, and the smallest is said to be the floor of Φ.[3] We now define the *value-range* of a

[3]Strictly speaking, the supremum of a set S of numbers may be the smallest number that is *higher* than any number in S, and analogously for infimum. However, in view of Theorem

formula Φ as the interval $[\mathbf{f}(\Phi), \mathbf{c}(\Phi)]$. The value-range of an atomic proposition like Pa, for example, is [0,1], because there are models in which a is a full member of P, and there are models in which a is a full non-member of P. But the value-range of a formula of the form $\phi \vee \neg\phi$ will be $[0.5, 1]$, since no model can give $\phi \vee \neg\phi$ a value less than 0.5. Similarly, the value-range of $\phi \wedge \neg\phi$ is $[0, 0.5]$.

The new operators $\owedge$ and $\ovee$ have the same syntax as $\wedge$ and $\vee$, but a different semantics. Informally, the semantics can be described as tracking the truth-value of the formula relative to its value-range: given a certain such range, the semantics measures how much truer the formula is than its floor (the falsest it can get), and divides that by its full value-range. The semantics is formalized as follows:

(8)
$$v(\phi \owedge \psi) = \begin{cases} v(\phi \wedge \psi) & \text{if } \mathbf{c}(\phi \wedge \psi) = \mathbf{f}(\phi \wedge \psi) \\ \frac{v(\phi\wedge\psi) - \mathbf{f}(\phi\wedge\psi)}{\mathbf{c}(\phi\wedge\psi) - \mathbf{f}(\phi\wedge\psi)} & \text{otherwise} \end{cases}$$

(9)
$$v(\phi \ovee \psi) = \begin{cases} v(\phi \vee \psi) & \text{if } \mathbf{c}(\phi \vee \psi) = \mathbf{f}(\phi \vee \psi) \\ \frac{v(\phi\vee\psi) - \mathbf{f}(\phi\vee\psi)}{\mathbf{c}(\phi\vee\psi) - \mathbf{f}(\phi\vee\psi)} & \text{otherwise} \end{cases}$$

As already mentioned, the main difference between $\owedge$ and $\ovee$ and their truth-functional counterparts $\wedge$ and $\vee$ is that $\owedge$ and $\ovee$ make sure that the range of possible truth values of their result includes both 0 and 1 whenever this is possible. Both functions, $\owedge$ and $\ovee$, accomplish this by scaling up the range of possible truth values. Such scaling is impossible only if the range of possible truth values of a formula is only a single point and in this case, the value remains the single dot. On a classical view, the range of truth-values for contradictions like $Pa \wedge \neg Pa$ and tautologies like $Pa \vee \neg Pa$ is only a single dot. But on a multi-valued approach, the intervals can have non-zero length: [0,0.5] for borderline contradictions and [0.5,1] for corresponding tautologies. So, the rescaling has an effect in multi-valued semantics.[4]

2, since for any formula ϕ and any number i in [0,1], ϕ takes i in some model, we need not worry that the sup and inf are values not actually taken.

[4]Our semantics is recursive, although not compositional on v_M. By switching to a more abstract semantic value, a function from models to valuations, we do get a compositional semantics.

Let $\mu(\phi$ be the function from the set of models $\mathcal{M}$ into $[0,1]$ with $\mu(\phi)(M) = v_M(\phi))$.
Let $\mathbf{c}(\mu(\phi)) = \sup\{k : \text{for some model } M, \mu(\phi)(M) = k\}$.
Let $\mathbf{f}(\mu(\phi)) = \inf\{k : \text{for some model } M, \mu(\phi)(M) = k\}$.

In our illustration of the system, we hereafter focus on models where the atomic propositions of the language – i.e. simple predications of the form Pa, Qb, etc. – are in their totality fully vague and independent as defined as follows:

> **Definition:** A set of models $\mathcal{M}$ for language L is *fully vague and independent* if and only if for any set of atomic propositions $\{p_1, \cdots, p_n\}$, and for any set of truth values $\{t_1, \cdots, t_n\} \subseteq [0,1]$, that there is a model M such that $v_M(p_1) = t_1, v_M(p_2), \ldots, v_M(p_n) = t_n$.

For linguistic applications, the full vagueness and independence assumption likely only hold for a subset of all atomic propositions: Concepts like *even number* may not be vague and concepts like *tall* and *giant* may not be independent of each other. Our system can, of course, be applied in such linguistically realistic models as well. However, the system's logical properties are likely to differ from the ones we explore below outside the subset of the language that only consists of fully vague and independent atomic propositions and their logical combinations.

An interesting logical property of fully vague and independent models is illustrated by the following graph. We consider here the case of just one atomic formula a with models $\mathcal{M} = [0,1]$ where $V_M(a) = M$ for any $M \in \mathcal{M}$. As the graph illustrates, for ϕ from $a \owedge \neg a$, $(a \owedge \neg a) \owedge a$, and $(a \owedge \neg a) \owedge \neg(a \owedge \neg a)$ the set of M with $V_M(\phi) = x$ can differ: e.g. $V_M(a \owedge \neg a) = 0$ has two solutions, but $V_M((a \owedge \neg a) \owedge a) = 0$ has three. But the set of solutions of $V_M(\phi) = x$ remains always finite for any $x \in [0,1]$. This is a more general property: for any sentence ϕ, $V_M(\phi) = x$ can never have infinitely many solutions as we prove below. This entails then that no ϕ is constant and, therefore, that no sentence is always true in a fully vague and independent model.

Clearly, the values of the **c** and **f** functions are determined by their arguments alone. Then there is a function $r_{\owedge}$ such that $\mu(\phi \owedge \psi) = r_{\owedge}(\mu(\phi), \mu(\psi))$. Define $r_{\owedge}$ such that

$$r_{\owedge}(m, m')(M) = \frac{\min(m(M), m'(M)) - \mathbf{f}(\min(m, m'))}{\mathbf{c}(\min(m, m')) - \mathbf{f}(\min(m, m'))}$$

where m, m' are meanings (values of μ). Similarly for $\ovee$.

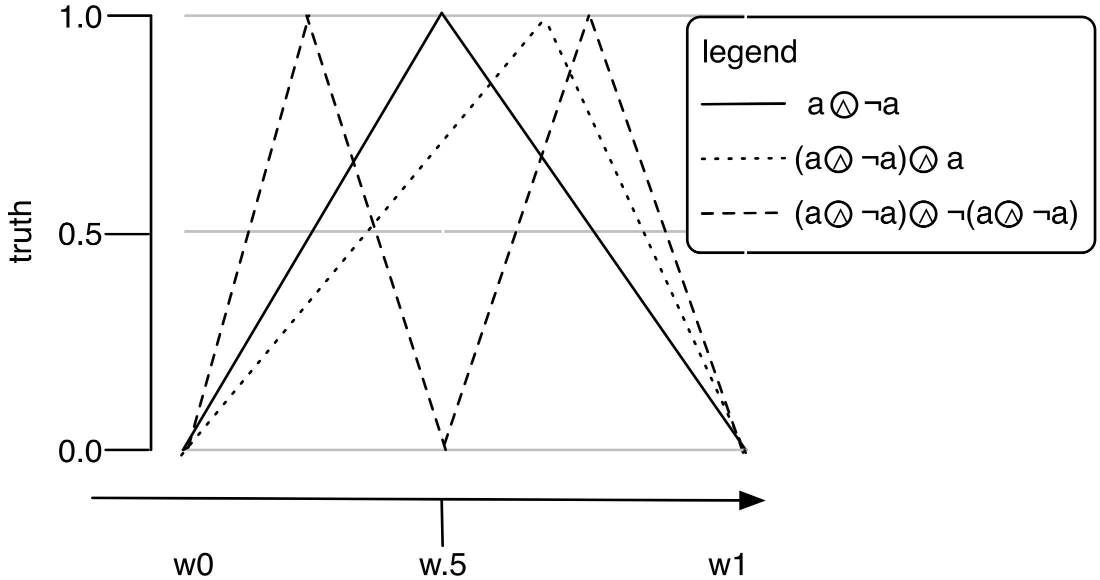

The following is the central theorem for fully vague and independent models: In the definitions of $\ovee$ and $\owedge$, the special case of the floor and ceiling being identical is covered. The following theorem shows that this special case is unnecessary for fully vague and independent models – the possible values of any formula in the system are always some non-trivial interval.

> **Theorem:** If $\mathcal{M}$ is a fully vague and independent set of models for language L, then for any sentence ϕ, the set $\{V_M(\phi) \mid M \in \mathcal{M}\}$ is equal to an interval $[a, b]$ with $a < b$.
>
> **Proof**: Assume that there are n atomic formulae, $a_1, \ldots, a_n$ that occur in ϕ. The assumption that $\mathcal{M}$ is fully vague and independent entails, that, without loss of generality, we can assume that the set of models be $\mathcal{M} = [0, 1]^n$ and that $V_M(a_i) = M_i$ for all $M \in \mathcal{M}$ and $i \in \{1, \ldots, n\}$. Let μ be the Lebesgue-measure on $\mathcal{M}$ with $\mu(\mathcal{M}) = 1$.
>
> We now show by induction over the structure of the formula that for any well-formed formula ϕ there exist $a, b \in [0, 1]$ with $a < b$ such that the three conditions (i) to (iii) hold. From condition (i), the theorem follows straightforwardly.
>
> (i) $\forall x \in [a, b] \exists M \in \mathcal{M} : V_M(\phi) = x$

(ii) $\mu(\{M \in \mathcal{M} \mid V_M(\phi) = x\}) = 0$

(iii) $M \mapsto V_M(\phi)$ is a continuous function from $M = [0,1]^n$ to $[0,1]$

In the base case, ϕ is an atomic formula a_i and $\{M \in \mathcal{M} \mid V_M(\phi) = x\} = \{M \in \mathcal{M} \mid M_i = x\}$, so (i), (ii), and (iii) evidently hold.

If ϕ is a complex sentence, either $\phi = \neg\alpha$ or it's made up of α and β conjoined by one of the four binary operators $\ovee$, $\owedge$, $\vee$, and $\wedge$. If $\phi = \neg\alpha$, then $\{M \in \mathcal{M} \mid V_M(\phi) = x\} = \{M \in \mathcal{M} \mid V_M(\alpha) = 1 - x\}$, and therefore (i) and (ii) hold by the inductive hypothesis. Similarly (iii) follows from the induction hypothesis because $V_M(\phi) = 1 - V_M(\alpha)$. Among the binary operators, $\ovee$ can be treated analogously to $\owedge$ and $\vee$ analogously to $\wedge$. We consider only the cases of $\wedge$ and $\owedge$ in detail.

Condition (iii) holds for $\wedge$ because the function $M \mapsto V_M(\alpha \wedge \beta)$ has as its value $\min(V_M(\alpha), V_M(\beta))$ and the minimum of two functions that are continuous is also continuous.

Condition (ii) for $\wedge$ follows because for any $y \in [0,1]$ the following holds by induction hypothesis and basic assumptions of measure theory:

$$\begin{aligned}
& \mu(\{M \in \mathcal{M} \mid V_M(\alpha \wedge \beta) = y\}) \\
= \; & \mu(\{M \in \mathcal{M} \mid V_M(\alpha) = y \;\&\; V_M(\beta) > y\}) \\
& + \mu(\{M \in \mathcal{M} \mid V_M(\alpha) \geq y \;\&\; V_M(\beta) = y\}) \\
\leq \; & \mu(\{M \in \mathcal{M} \mid V_M(\alpha) = y\}) + \mu(\{M \in \mathcal{M} \mid V_M(\beta) = y\}) \\
= \; & 0
\end{aligned}$$

Now set $a = \inf\{x \in [0,1] \mid \exists M \in \mathcal{M}\ V_M(\alpha \wedge \beta) = x\}$ and $b = \sup\{x \in [0,1] \mid \exists M \in \mathcal{M}\ V_M(\alpha \wedge \beta) = x\}$. If $a = b$, then the set $\{M \in \mathcal{M} \mid V_M(\alpha \wedge \beta) = a\}$ would be $\mathcal{M}$, but this cannot be because we just saw that $\mu(\{M \in \mathcal{M} \mid V_M(\alpha \vee \beta) = a\}) = 0$, while $\mu(\mathcal{M}) = 1$. Hence $a < b$ must hold.

Because $\mathcal{M}$ is a closed set and $M \mapsto V_M(\alpha \wedge \beta)$ is continuos in M, infimum and supremum of $M \mapsto V_M(\alpha \wedge \beta)$ must be actually attained – i.e. there must be $M_a, M_b \in \mathcal{M}$ with $V_{M_a}(\alpha \wedge \beta) = a$ and $V_{M_b}(\alpha \wedge \beta) = b$. Furthermore, for any $x \in (a,b)$ there must also be at least one $M \in \mathcal{M}$ with $V_M(\alpha \wedge \beta) = x$ because of the continuity of the function $M \mapsto V_M(\alpha \wedge \beta)$. This establishes that condition (i) holds for $\alpha \wedge \beta$.

Finally from the result for $\wedge$ it follows that (i), (ii), and (iii) also hold for $\phi = \alpha \owedge \beta$. Namely $V_M(\alpha \owedge \beta)$ is a linear function of $V_M(\alpha \wedge \beta)$ because the image of $M \mapsto V_M(\alpha \wedge \beta)$ is an interval $[a, b]$ with $b > a$. $\square$

Corollary: If $\mathcal{M}$ is fully vague and independent, any formula ϕ where the highest level binary operator is $\ovee$ or $\owedge$, has [0,1] as its image under $V_M\phi$ for $\mathcal{M}$. *Proof*: The unary operator $\neg$ can be ignored since $\{V_M(\neg\psi) \mid M \in \mathcal{M}\} = [0,1]$ if $\{V_M(\psi) \mid M \in \mathcal{M}\} = [0,1]$. Hence, we can assume $\phi = \alpha \ovee \beta$ or $\phi = \alpha \owedge \beta$. But the theorem entails that there are $a < b$ such that $\{V_M(\alpha \vee \beta) \mid M \in \mathcal{M}\} = [a, b]$ or $\{V_M(\alpha \wedge \beta) \mid M \in \mathcal{M}\} = [a, b]$. From the definition of $\ovee$ and $\owedge$ the corollary then follows. $\square$

Now, let a *logical truth* be a sentence that under any valuation takes only *designated* truth values. Any *partition* of the set of truth values must consist of a non-empty set D of *designated values* and a non-empty set N of non-designated truth values. D and N must be mutually exclusive and jointly exhaustive.

Corollary: If $\mathcal{M}$ is fully vague and independent, there are no logical truths in L where the highest level, binary operator is $\ovee$ or $\owedge$. *Proof*: Since every sentence where the highest level, binary operator is $\ovee$ or $\owedge$ takes every truth value, no sentence takes only designated values, whatever they are. $\square$

The predictions of the system we propose are similar to those of CERvR: we predict not just the acceptability of borderline contradictions, but also the unacceptability of the conjuncts that make them. The reason is that while contradictions are evaluated in the range $[0, 0.5]$, the conjuncts that comprise them are each evaluated in the full range of truth-values. When $v(Pa) = 0.5$, in borderline cases that is, Pa and $\neg Pa$ will each be half-true relative to the range [0,1]. Suppose now that some designated value d divides acceptable expressions from unacceptable ones, i.e. Φ is acceptable only if $v(\Phi) \geq d$. If d is set at any value greater than 0.5, say 0.7, then neither Pa nor $\neg Pa$ will be acceptable when a is a borderline case, because each is mapped to 0.5 by the valuation function v, and 0.5 < d. But when they are conjoined together by the operator $\owedge$, the resulting formula will be mapped to 1 by v:

$$v(Pa \owedge \neg Pa) = \frac{v(Pa \wedge \neg Pa) - \mathbf{f}(Pa \wedge \neg Pa)}{\mathbf{c}(Pa \wedge \neg Pa) - \mathbf{f}(Pa \wedge \neg Pa)}$$
$$= \frac{0.5 - 0}{0.5 - 0}$$
$$= 1$$

The proposal also differs from the multi-valued logic presented above on the disjunction of a borderline proposition with its negation. Serchuk et al. (2010) show that, empirically, classical tautologies like (10) are not always judged true. Specifically, if Joe is of borderline height, (10) is on average judged to be fully true. The multivalued logic we are considering predicts that the truth value of $Pa \vee \neg Pa$ could be minimally 0.5 and maximally 1.0. In actual fact, though, (10) is far less acceptable than its negation if Joe is of borderline height. This is not predicted by the definition of $\vee$. But, with the new disjunction $\ovee$ the expression $Pa \ovee \neg Pa$ can have truth value 0, because $\ovee$ is defined so as to scale the [0.5,1] range of $Pa \vee \neg Pa$ to the [0,1] interval.

(10) Joe is tall or not tall.

We now show how our account handles the cases that we think are problematic for CERvR, namely the cases where contradictions and tautologies are embedded in larger expressions, e.g. our example (5), repeated here as (11)

(11) Every blond man is either rich or tall and not tall.

To simplify the discussion, let us put aside the quantificational element in (11) and focus on the proposition $(\phi \owedge \neg\phi) \ovee \psi$[5]. Let us also assume that $v(\phi) = 0.5$ and $v(\psi) = 0$ (Joe is of borderline height, and he is poor). On CERvR's account, this expression must be evaluated strictly, because there are models in which the logically independent proposition ψ is strictly true. But since ψ happens to be false in M, the truth of the expression will rest on the truth of $\phi \wedge \neg\phi$, and since the evaluation of truth is strict, the entire expression will come out false. On the system presented here, the acceptability of the

[5]A more suitable English example of this proposition would be "John is either rich or tall and not tall". But since the cited experiments used sentences like (11), we chose to use quantified sentences as examples here, but to ignore the quantifier in our discussion of the logic. We hope that future experimental work will show whether non-quantificational sentences behave in the same way.

expression will equal the acceptability of the contradiction, which will in turn be acceptable when ϕ is a borderline case:

$$
\begin{aligned}
v((\phi \owedge \neg\phi) \ovee \psi) &= \frac{v((\phi \owedge \neg\phi) \vee \psi) \, - \, \mathbf{f}((\phi \owedge \neg\phi) \vee \psi)}{\mathbf{c}((\phi \owedge \neg\phi) \vee \psi) \, - \, \mathbf{f}((\phi \owedge \neg\phi) \vee \psi)} \\
&= \frac{v((\phi \owedge \neg\phi) \vee \psi) \, - \, 0}{1 \, - \, 0} \\
&= v((\phi \owedge \neg\phi) \vee \psi) \\
&= v(\phi \owedge \neg\phi) \qquad\qquad \text{(since } v(\psi) = 0\text{)}
\end{aligned}
$$

A few other predictions are worth noting: we no longer assign $\phi \owedge \phi$ and $\phi \owedge \neg\phi$ the same status in borderline cases. Kamp (1975), Fine (1975) and others have criticized fuzzy logic for assigning the same truth-value to these conjunctions in the borderline range – when $v(\phi) = 0.5$, both conjuncts will have the value 0.5. The proposed approach makes different predictions: $\phi \owedge \phi$ will be acceptable whenever ϕ itself is acceptable because adding ϕ in a conjunction will have no effect on the value-range of the formula, but conjoining ϕ with its negation will shrink the value-range to [0,0.5], and as was already demonstrated, this makes the contradiction more acceptable than any of ϕ, $\neg\phi$, $\phi \owedge \phi$, and $\neg\phi \owedge \neg\phi$. We also predict a difference between the acceptability of $Pa \owedge \neg Pa$ on the one hand, and $Pa \owedge \neg Qa$ on the other. In fuzzy logic, these two formulae are assigned the same truth-value whenever Pa and Qa are both half-true. But if Pa and Qa are logically-independent, the value-range of the second formula will be [0,1], rather than [0,0.5]. So, given a designated value $d > 0.5$, we predict $Pa \owedge \neg Qa$ to be unacceptable, but we predict the contradiction $Pa \owedge \neg Pa$ to be acceptable. This result was tested and verified empirically by Sauerland (2011).[6]

There is another reason for keeping the designated value d above 0.5. Consider a case where $v_M(Pa) = 0.3$ and $v_M(Pb) = 0.5$. If d was 0.3, we would predict that the sentence $Pa \owedge \neg Pb$ be assertable, because it has value 0.3. This, however, is undesirable because a has property P to a lesser degree than b. Generally, this problematic set-up can only arise if $d \leq 0.5$; the set-up requires that

[6]From a computational point of view, the proposed semantics looks rather complex. But since it, by current empirical findings, models usage well, verdicts about computational complexity should best be suspended for the time being.

(i) $v(Pa \owedge \neg Pb) \geq d$, and

(ii) $v(Pa) \leq v(Pb)$.

From (i) we get $v(Pa) \geq d$ and $v(\neg Pb) \geq d$, that is

(iii) $v(Pa) \geq d$

(iv) $1 - v(Pb) \geq d$

Combining (ii) and (iv) produces $1 - v(Pa) \geq d$, which, when added to (iii) gives use

(v) $v(Pa) + (1 - v(Pa)) \geq 2d$

$1 \geq 2d$

$0.5 \geq d$

It is therefore necessary for d to be no greater than 0.5 if this undesirable outcome were to arise. Setting d strictly above 0.5 will overcome it, since it will make the conjunction unassertable.

Finally, we can define a natural notion of logical consequence as follows: Let ψ be a *logical consequence* of Γ iff in every model M there is at least one sentence $\phi \in \Gamma$ such that $v_M(\phi) \leq v_M(\psi)$.

Logical consequence in fully vague and independent models is almost as thin as logical truth, but not quite. Negation is classical, and hence we have $\phi \models \neg\neg\phi$ as well as $\neg\neg\phi \models \phi$.

Because $\phi \owedge \neg\phi$ has value 1 when ϕ has value 0.5, it has a higher value than both ϕ and $\neg\phi$, and so $\phi \owedge \psi \not\models \phi$ and $\phi \owedge \psi \not\models \psi$ – i.e. conjunction elimination is not valid. Similarly, since $\phi \ovee \neg\phi$ has value 0 when ϕ has value 0.5, it has a lower value than both ϕ and $\neg\phi$, and so $\phi \not\models \phi \ovee \psi$ and $\psi \not\models \phi \ovee \psi$.

The *ex falso quodlibet* consequence $\phi, \neg\phi \models \psi$, doesn't hold in fully vague and independent models. But, notice that because of the failure of conjunction elimination, we could add the rule $\phi, \neg\phi \models \psi$, without inconsistency. Hence, the logic of L is *not* paraconsistent.

Note, finally, that although conjunction introduction and disjunction elimination both fail, the other halves, conjunction introduction and disjunction elimination hold.

Theorem: The following hold in L if the highest level, binary operator of ϕ, ψ, and ξ is $\ovee$ or $\owedge$:

a) $\phi, \psi \models \phi \owedge \psi$
b) If $\phi \models \xi$ and $\psi \models \xi$, then $\phi \ovee \psi \models \xi$.

Proof: Consider first a). Let $M \in \mathcal{M}$ be arbitrary. The first corollary above entails $c(\phi \wedge \psi) = 0$. Hence $v_M(\phi \owedge \psi) = \frac{v_M(\phi \wedge \psi)}{c(\phi \wedge \psi)}$, which is greater than or equal to $v_M(\phi \wedge \psi)$ because $c(\phi \wedge \psi) \in (0, 1])$. Because $v_M(\phi \wedge \psi) = \min(v_M(\phi), v_M(\psi))$, it follows that $v_M(\phi \owedge \psi) \geq v_M(\phi)$ or $v_M(\phi \owedge \psi) \geq v_M(\psi)$. The proof of b) is analogous. $\square$

4 Conclusion

We have shown that examples like (12) (= 4 above) are problematic for the approach to borderline contradictions based on linguistic pragmatics of CERvR. Their proposal predicts (12) to be false even in a situation where its first disjunct *Joe is and isn't tall* is judged true.

(12) Joe is and isn't tall or Joe has red hair.

To remedy the problem, we developed a semantic proposal to account for borderline contradictions. The proposal is based on new, intensional entries for disjunction as $\ovee$ and conjunction as $\owedge$.

Our proposal predicts an interesting difference between formulas like $(\phi \owedge \neg\phi) \ovee \psi$ and $(\phi \ovee \psi) \owedge \neg\phi$, where ϕ is a borderline statement and ψ is logically independent of ϕ. As we showed above, our proposal predicts that $(\phi \owedge \neg\phi) \ovee \psi$ be judged true. However, it doesn't make the same prediction for $(\phi \ovee \psi) \owedge \neg\phi$: because $\phi \vee \psi$ has the full range [0,1] of potential truth values, and therefore $\phi \ovee \psi$ is predicted to have the same truth value as $\phi \vee \psi$. But, then $(\phi \ovee \psi) \ovee \neg\phi$ also has the full range [0,1] of possible truth values, and therefore it means the same as $(\phi \ovee \psi) \owedge \neg\phi$. We found it difficult to come up with natural language examples to test this prediction. Sentences like *Joe is tall or has red hair and he is not tall* are so awkward that it does not seem possible to pass reliable judgments on them.

Bibliography

Alxatib, Sam & Jeffrey Pelletier. 2010. On the psychology of truth-gaps. In Rick Nouwen, Robert van Rooij, Uli Sauerland & Hans-Christian Schmitz (eds.), *Vagueness in communication*. Heidelberg, Germany: Springer. doi: 10.1007/978-3-642-18446-8.

Bonini, Nicolao, Daniel Osherson, Riccardo Viale & Timothy Williamson. 1999. On the psychology of vague predicates. *Mind & Language* 14. 377–393.

Cobreros, P., P. Egré, D. Ripley & R. van Rooij. 2011. Tolerant, classical, strict. *Journal of Philosophical Logic* .

Dalrymple, Mary, Makoto Kanazawa, Yookyung Kim, Sam Mchombo & Stanley Peters. 1998. Reciprocal expressions and the concept of reciprocity. *Linguistics and Philosophy* 21. 159–210.

Fine, Kit. 1975. Vagueness, truth and logic. *Synthese* 30. 265–300.

Hajek, Petr. 2009. Fuzzy logic. In Edward N. Zalta (ed.), *The stanford encyclopedia of philosophy*. Spring 2009 edn.

Kamp, Hans. 1975. Two theories about adjectives. In Edward L. Keenan (ed.), *Formal semantics of natural language*. Cambridge, England: Cambridge University Press.

Kamp, Hans & Barbara H. Partee. 1995. Prototype theory and compositionality. *Cognition* 57. 129–191.

Ripley, David. 2011. Contradiction at the borders. In Rick Nouwen, Robert van Rooij, Uli Sauerland & Hans Christian Schmitz (eds.), *Vagueness in communication*, 175–94. Heidelberg: Springer.

Sauerland, Uli. 2011. Vagueness in language: The case against fuzzy logic revisited.

Sauerland, Uli, Peter Pagin, Sam Alxatib & Stephanie Solt. 2011. Vagueness and the semantics and pragmatics of contradiction. Presentation Delivered at the Euro XPRAG Workshop in Pisa, Italy.

Serchuk, Phil, Ian Hargreaves & Richard Zach. 2010. Vagueness, logic and use: Four experimental studies on vagueness.