Are Gaps Preferred to Gluts? A Closer Look at Borderline Contradictions

Paul Égré and Jérémy Zehr

Abstract This paper examines the acceptance of so-called borderline contradictions involving vague adjectives. A close look at the available data from previous studies points toward a preference for "gappy" descriptions of the form "*x* is neither *P* nor not *P*" over "glutty" descriptions of the form "*x* is *P* and not *P*". We present the results of an experiment in which we tested for that difference systematically, using relative gradable adjectives. Our findings confirm that both kinds of descriptions are accepted, but indeed that "neither"-descriptions are to a large extent preferred to "and"-descriptions. We examine several possible explanations for that preference. Our account relies on the distinction proposed by Cobreros et al. (J Philos Logic, 1–39, 2012) between *strict* and *tolerant* meaning for vague adjectives, as well as on a specific implementation of the strongest meaning hypothesis endorsed by Cobreros et al. as well as Alxatib and Pelletier (Mind Lang 26(3): 287–326 2011a). Our approach, however, argues in favor of local pragmatic strengthening instead of global strengthening in order to derive that preference.

Keywords Borderline contradictions · Strict-tolerant logic · Gradable adjectives Vagueness · Presupposition · Local accommodation · Pragmatic strengthening Three-valued logic

P. Égré (✉)
Département d'Études Cognitives et Département de Philosophie de l'École Normale Supérieure, Institut Jean-Nicod, ENS, CNRS, EHESS, PSL Research University, Paris, France
e-mail: paul.egre@ens.fr

P. Égré
Swedish Collegium for Advanced Study, Uppsala, Sweden

J. Zehr
University of Pennsylvania, Philadelphia, PA 19104, USA
e-mail: jeremy.e.zehr@gmail.com

E. Castroviejo et al. (eds.), *The Semantics of Gradability, Vagueness, and Scale Structure*, Language, Cognition, and Mind 4, https://doi.org/10.1007/978-3-319-77791-7_2

1 Introduction

Several experiments in the last decade indicate that sentences that correspond to literal contradictions from the standpoint of classical logic are accepted to a significant extent by naive subjects to describe *borderline cases* of a vague predicate, in a way that they are not for so-called *clear cases* (Ripley 2011a; Alxatib and Pelletier 2011a; Serchuk et al. 2011; Égré et al. 2013). Thus, Ripley (2011a) found that sentences like (1-a) and (1-b) are accepted significantly more by subjects presented with a circle-square pair at middling distance from one another, compared to a circle and a square touching each other, or to a circle and a square appearing further away from each other. Similarly, Alxatib and Pelletier (2011a) found that sentences of the form (2-a) and (2-b) are checked "True" significantly more by participants when pertaining to a man appearing of height 5′11″ than for a man appearing of height 6′6″ or for a man of height 5′4″.

(1) a. The circle is and isn't near the square.
b. The circle neither is nor isn't near the square. Ripley (2011a)

(2) a. Man x is tall and not tall.
b. Man x is neither tall nor not tall. Alxatib and Pelletier (2011a)

Similar sentence forms are intuitively unacceptable when involving *precise* predicates (see (3)), setting aside cases of presupposition failure as in (4). Possibly, (4-b) might be acceptable, but only to convey that $\sqrt{2}$ is outside the domain of application of the predicate "prime number", assuming the latter applies only to integers (see Zehr 2014 for more on the difference between vagueness and presupposition failure)[1]

(3) a. # 9 is and isn't a prime number.
b. # 9 neither is nor isn't a prime number.

(4) a. # $\sqrt{2}$ is and isn't a prime number.
b. (?) $\sqrt{2}$ neither is nor isn't a prime number.

Prima facie, the acceptability of contradictory sentences in borderline cases (so-called "borderline contradictions", following Ripley 2011a's terminology) may not appear so surprising, since a borderline case of a vague predicate is often characterized as a case for which one feels equally attracted toward applying and toward denying the predicate (e.g. Peirce 1902). From a behavioral point of view, however,

[1] Sentence (4)-b may be judged outright false of course, since $\sqrt{2}$ *is not* a prime number. It seems to us acceptable in a context in which a teacher, let us say, would want to cut short a dispute between two pupils, one of them arguing that $\sqrt{2}$ is a prime integer, and the other arguing that $\sqrt{2}$ is not a prime integer, both mistakenly thinking it is an integer. Both pupils would wrongly presuppose that $\sqrt{2}$ is an integer, and the point of (4)-b would be to reject that presupposition. See Sect. 4.4.2 below for more on the analysis of such presuppositional sentences.

that characterization is compatible with subjects systematically rejecting descriptions of the form "x is P and not P", or "x is neither P nor not P", while judging "x is P" half the time, and denying "x is P" the other half.

The results of Ripley (2011a); Alxatib and Pelletier (2011a) suggest that that picture is inadequate, however. From a semantic point of view, those findings are not easily accommodated by either supervaluationist or subvaluationist theories of vagueness, which predict sentences of the form "x is P and not P" and "x is neither P nor not P" to remain contradictions in borderline cases (see Ripley 2013; Alxatib and Pelletier 2011a; Égré et al. 2013). They are more easily accommodated in paraconsistent-friendly frameworks or in fuzzy logic, however, that is in theories in which classical contradictions do not automatically receive the value False. One framework of particular to us is the so-called strict-tolerant framework (Cobreros et al. 2012), in which both kinds of sentences can be true tolerantly in borderline cases.

In this paper, we propose to dig further into the explanation of borderline contradictions. In what follows we shall refer to sentences of the form "x is P and not P" as *conjunctive* descriptions of borderline cases, or *conjunctions* for short, and to sentences of the form "x is neither P nor not P" as *negative disjunctive* descriptions, or *disjunctions* for short. More often, we will refer to them simply as "and"-descriptions and "neither"-descriptions. Although we shall sometimes use the expressions "glutty" descriptions and "gappy" descriptions, we will favor the "and" versus "neither" denominations, which in a sense are the most theory-neutral. The phenomenon we are interested in is whether the two kinds of descriptions are equally accepted in borderline cases, or whether there is a preference for one kind of description over the other. Our intuition tells us that disjunctions might be preferred to conjunctions, that is, it might be easier to describe a borderline case as "neither tall nor not tall" than as "tall and not tall".

A closer look at extant results suggests that this is likely to be the case (see the review in Sect. 2). The difference, however, has not been an object of attention in previous studies. We propose to test for that preference. One motivation to do so is that the difference can potentially cast further light on the selection between two kinds of meaning for vague predicates. Alxatib and Pelletier (2011a, b) and Cobreros et al. (2012, 2015b) both treat vague predicates as pragmatically ambiguous between a strong and a weak meaning (aka *strict* and *tolerant* meanings, see Cobreros et al. 2012). The strong or strict meaning of a predicate like "tall" is intuitively equivalent to "clearly tall" and the weak or tolerant meaning to "relatively tall" ("not clearly not tall"). Alxatib and Pelletier (2011a, b) and Cobreros et al. (2012, 2015b) both formulate the hypothesis that the strong meaning ought to be selected first, but do not look at whether the principle should entail a preference for negated disjunctions over conjunctions. In Sect. 3, we present the results of an experiment confirming our main intuition. In Sect. 4, we propose algorithm intended to account for that result. This algorithm too implements the idea that the strict meaning is selected before the tolerant meaning, but importantly it rests on the idea that pragmatic strengthening is done locally, rather than globally for whole sentences.

2 Gaps and Gluts: A Brief Review

Borderline cases of a vague predicate are commonly described either as cases leaving a gap between the positive extension of the predicate and its negative extension, or in a dual manner, as cases where the positive extension of the predicate and its negative extension overlap, thereby creating a glut (see Fine 1975; Égré et al. 2013). Intuitively, a negative disjunctive description of the form "x is neither P nor not P" matches the idea of a gap between the positive and the negative extension, and a conjunctive description of the form "x is both P and not P" the idea of a glut. The first question we seek to clarify is whether "glutty" and "gappy" descriptions of borderline cases are used to the same extent.

Ripley (2011a) presented participants with two kinds of conjunctive and disjunctive descriptions, which he calls elided versus non-elided. In an elided conjunction, the conjunction is internal to the VP ("The circle both is and isn't near the square"), whereas in a non-elided conjunction, it is sentential ("The circle is near the square and it isn't near the square"). In both what Ripley calls non-elided disjunctions ("The circle neither is near the square nor isn't near the square") and elided disjunctions ("The circle neither is nor isn't near the square"), the disjunction is VP-internal, but what varies is whether part of the VP is elided or not. Ripley presented participants with 7 pairs of a circle and square at varying distances from one another, and asked them to rate each description for each stimulus on a 1–7 scale, with 1 labeled 'Disagree' and 7 labeled 'Agree'. In Fig. 1, we give a summary of his data, where we aggregated scores for so-called elided versus non-elided description types. Prima facie, we see no preference for one description type over the other, neither globally, nor in the specific case of stimulus C which receives the highest assent to both description types.

On the other hand, we do observe an overall preference for disjunctions over conjunctions in the experiment run by Alxatib and Pelletier (2011a). Alxatib and Pelletier's methodology was different from Ripley's, since participants had to check True, False, or Can't Tell to four kinds of description including "Tall", "not Tall", "Tall

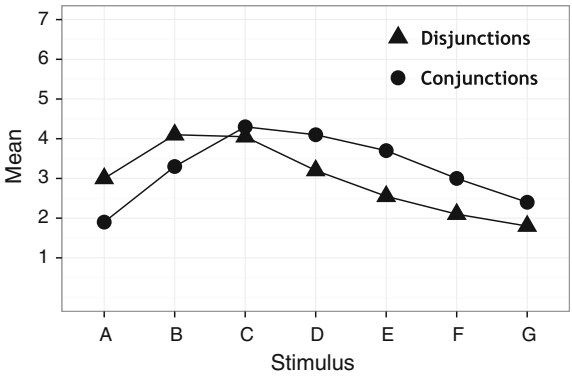

Fig. 1 Ripley 2011's data: triangles represent aggregate scores for disjunctions, and dots aggregate scores for conjunctions. Mean scores are highest for non-extreme distances (B and C)

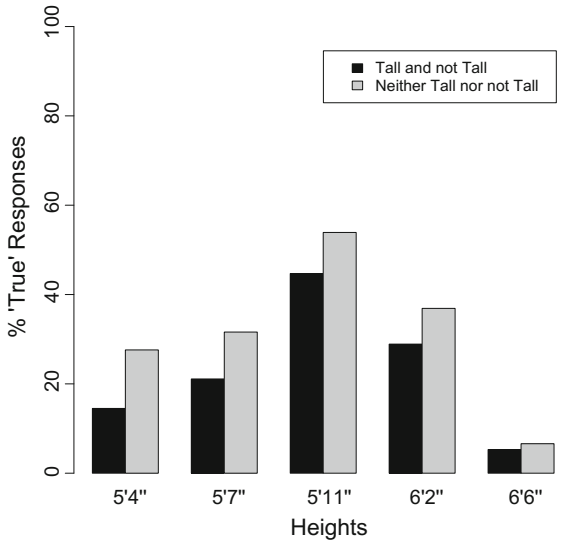

Fig. 2 Alxatib and Pelletier (2011a)'s data: proportion of 'True' checks in conjunctions and disjunctions

Table 1 Alxatib and Pelletier (2011a)'s data for the central stimulus of height $5'11''$ (n and b refer to "neither"- vs. "both"-sentences, T, F, C to "True", "False" and "Can't tell" responses)

	Tn	Fn	Cn	Total
Tb	22	12	0	34
Fb	13	18	0	31
Cb	6	2	3	11
Total	41	32	3	76

and not Tall", and "Neither tall nor not tall". In Fig. 2, we reproduce the proportions of 'True' checks to the latter two kinds, namely conjunctions and disjunctions. What we can observe is a higher level of 'True' checks to disjunctions, in each stimulus. For the stimulus of middling height in their stimulus set, Alxatib and Pelletier (2011b) report the data reproduced in Table 1, where Tn, Fn, and Cn stand for the numbers of participants responding "True", "False" and "Can't tell" to "#2 is neither tall nor not tall", and where Tb, Fb, and Cb give the corresponding numbers for "#2 is both tall and not tall". As they point out, a McNemar-Bowker test for symmetry gives a value of 8.04, with $p < 0.05$, suggestive of a difference between the two descriptions.

Serchuk and colleagues also investigated borderline contradictions, this time comparing more explicitly than Ripley the effect of having two kinds of negation, VP-internal or sentential. Unlike Ripley or Alxatib and Pelletier, they did not use perceptual stimuli, but asked participants to imagine borderline cases, describing those in terms of the operator "clearly". For instance, in one of their scenarios, a woman named Susan is described as being "somewhere between women who are clearly rich and women who are clearly non-rich". For disjunctions, they used *unnegated*

Table 2 Serchuk et al. (2011)'s data on disjunctive and conjunctive descriptions of borderline cases

Description	'True'	'False'	Other
Either x is P or x is not P	113	137	100
x is P or it is not the case that x is P	141	88	121
x is P and x is not P	66	195	89
x is P and it is not the case that x is P	25	247	78

disjunctions of two kinds: "Either Susan is rich or Susan is not rich", and "Susan is rich or it is not the case that Susan is rich". Each time, however, the disjunction is sentential and not VP-internal: they did not present participants with "Susan is rich or not rich". They used a similar pair of probes for conjunctions, namely "Susan is rich and Susan is not rich", and "Susan is rich and it is not the case that Susan is rich" (here too, they did not use "Susan is rich and not rich"). Participants had to check exactly one answer among six possible answers for each sentence, within the set consisting of {"true", "not true but also not false", "partially true and partially false", "false", "both true and false", "true or false but I don't know which"}. In Table 2, we present Serchuk et al.'s data in a table in which we collapsed the responses other than True and False under a third category "Other". In this case, we cannot directly compare "True" responses to "neither" and "and" descriptions, because Serchuk et al. did not use any sentence of the form "neither... nor". However, we can get an indirect idea of the participants' judgments about "neither" sentences, granted that they would stand in an inverse relation to those for "either" sentences. We shall therefore compare the proportion of "True" responses for conjunctions to the proportion of "False" responses for disjunctions. There is a higher proportion of False answers to the disjunction "either x is P or x is not P" than of "True" answers to "x is P and not P", and similarly when we compare "False" answers to "x is P or it is not the case that x is P" with "True" answers to "x is P and it is not the case that x is P". A Fisher test based on that comparison yields a significant difference for the former two sentences ($p < 10^{-10}$) as well as for the latter two ($p < 10^{-16}$).

Finally, we mention the results of an unpublished study, by Solt and Gotzner (2010). They showed participants picture series depicting either suitcases of various sizes, bluejeans of varying price, or cities of different distances from Berlin. The goal of the experiment was to see which pictures in each series would be classified by participants as satisfying the respective adjectives "big" (*gross*), "expensive" (*teuer*), "far" (*weit*) in comparison to their negation or the polar antonym. In one condition participants had to decide which pictures satisfied the adjective, and which satisfied the syntactic negation. In another condition a different group had to decide which pictures satisfied the adjective, and which satisfied the polar antonym. Their results indicate that only a small minority of participants left no gap between the adjective and its opposite (whether polar or syntactic), or ascribed the same item to both

descriptions. Although that study does not investigate the acceptance of complex sentences like the ones we are interested in, it supports the idea of a preference for gaps over gluts when people are asked to judge of a sentence and its negation separately.

What the present review points toward, therefore, is a preference for the description of borderline cases in terms of negated disjunctions ('neither P nor not P') over conjunctions ('P and not P'). Because the data we reviewed are partial and unsystematic, we proceed to test that hypothesis more systematically in what follows. An important caveat is that we limited our study to sentences of the form "*x* is P and not P" and "*x* is neither P nor not P". We did not test sentences in which the conjunction or the disjunction are clearly sentential ("*x* is P and *x* is not P"/"neither is *x* P nor is *x* not P") nor sentences in which the negation is propositional ("it is not the case that"). From Serchuk et al.'s data, we see that there is an effect of those variations on acceptance, but what we were interested in is whether there is a contrast between conjunctive and disjunctive sentences for those cases in which the sentences seem best accepted in the first place.

3 Experiment

In order to test the preference for "gappy" descriptions over "glutty" descriptions of borderline cases, we designed an experiment intended to compare the acceptance of both kinds of descriptions. The experiment we report here is the third and main one in a series of three, and it replicates the results of the two previous, pilot versions. We presented participants with different scenarios involving different adjectives, each time involving the verbal description of a borderline case, drawing inspiration from Serchuk et al. (2011), who basically asked participants to imagine borderline cases based on similar verbal descriptions. We then asked participants to judge the adequacy of contradictory descriptions of the forms in (5), leaving them the possibility to either accept or reject the description.

(5) a. *Neither*-descriptions: "BORDERLINE is neither ADJECTIVE nor NOT ADJECTIVE"
b. *And*-descriptions: "BORDERLINE is ADJECTIVE and NOT ADJECTIVE"

Our prediction was that we should see a higher rate of acceptance of "neither"-descriptions over "and"-descriptions, but also that either should be significantly more accepted than outright false sentences. In order to get a homogenous set of results, however, our first step was to define general principles for the selection of adjectives and for the construction of our scenarios.

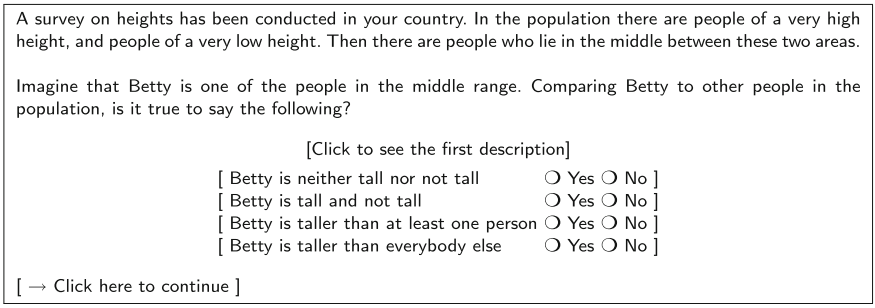

Fig. 3 Example of a trial. [Bracketed texts] appeared dynamically: each description appeared after answering the preceding one or clicking Click to see the first description. The set of answers was validated by clicking on Click here to continue (appearing after the last click)

3.1 Design

The experiment was a block design: each participant faced two series of descriptions. In one block, the descriptions were as exemplified in (5). In the other block, NOT ADJECTIVE was replaced with a corresponding antonym (e.g. *short* in place of *not tall*). In this paper, we do not report the results for these antonymic descriptions: we treat them as fillers, even though they served as critical conditions for the purposes of another study investigating the relation between the two types of negations (syntactic vs. lexical).

The design took the form of an acceptability task: participants were presented with a fictive scenario verbally depicting a borderline case on a given gradable dimension and had to tell whether they would accept four descriptions of this borderline case, as exemplified in Fig. 3. On each trial, they saw the scenario first, followed by the instruction *Click to see the first description*. The instruction would disappear after clicking and simultaneously reveal the first of the four descriptions to be assessed. Each time, clicking on one of *Yes* and *No* would dynamically let the next description appear on screen. For the last description, checking *Yes* or *No* would reveal a link needed to validate all four answers of the current trial, reading *Click here to continue*. As long as the descriptions were visible on the screen, the participants were able to modify their judgments. Once they clicked on *Click here to continue*, they would see the next trial, with a new scenario targeting a distinct adjective.

3.2 Materials

3.2.1 Scenarios

Serchuk et al. (2011) used verbal descriptions of borderline cases in their experiments, as exemplified in (6), and then asked their participants to judge a list of sentences about these borderline cases. What follows is an example of the sentences they used:

(6) "Imagine that on the spectrum of rich women, Susan is somewhere between women who are clearly rich and women who are clearly non-rich".

To the extent that their results revealed acceptance of contradictory descriptions, they show that participants readily represent borderline cases when explicitly asked to do so. This prompted us to describe borderline cases in similar verbal ways. However, in contrast with Serchuk et al.'s descriptions, we did not let the target adjectives (e.g. *rich* in (6)) appear in the scenarios, to avoid any priming effect. Indeed, the description in (6) might prime a *clearly* reading of *rich* that would lead to systematically exclude the borderline cases from the positive extension of *rich*. Alternatively, it might prime a contrastive, looser use of *rich* leading one to categorize the borderline cases in the positive extension. If any of these effects were real, they could bias our results in a way that does not appear to be the case in Alxatib and Pelletier (2011a)'s pictorial context. Because of that, we formulated our scenarios by referring to the *scale* associated with the adjectives, using morphologically unrelated nouns whenever possible (cf. Fig. 3 where we used the word *heights* to refer to the scale associated with *tall*). We hypothesized that participants would naturally represent the individuals in the middle of these scales as borderline cases for the target adjectives. Note that contrary to pictorial representations, verbal descriptions have the advantage of letting subjects build their own, ideal representations of what a borderline case might be on the given scale. The particular scenarios we used are reported in Appendix 1.

3.2.2 Selection of Adjectives: Four Principles

Not all adjectives come with borderline cases, and not all adjectives seem to define their borderline area in the same way. In order to have a set of adjectives as homogeneous as possible, we selected them along four criteria that we present here. Our selection was partly guided by principles investigated in the works of Roche (2012) and Ruytenbeek (2013), both of them conducted under the supervision of Benjamin Spector, concerning the negation of adjectives.[2]

(i) *Gradable*

As exemplified in the scenario in Fig. 3 all the borderline cases we described were presented as lying in the central region of *a given scalar dimension*.[3] All the adjectives

[2] See Ruytenbeek et al. (2017) for a published follow-up to that work.

[3] In this paper, we take this *non-extreme* property to be definitory of borderline cases. By contrast, one might consider that an *extremely* insane person in a *perfect* physical shape could be a borderline case

we used are gradable and therefore define a scale along which their arguments are non-trivially ordered. We say that an adjective is gradable if and only if at least one of the following constructions is perceived as natural.

(7) a. *X* is ADJECTIVE-er than *Y*.
b. *X* is more ADJECTIVE than *Y*.

For example *tall* is gradable (8-a), whereas *underage* is not (8-b).

(8) a. Bill is taller than Sue.
b. #Bill is more underage than Sue.

Criterion 1: We excluded any non-gradable adjective from our list.

(ii) *Relative Versus Absolute*

Two kinds of gradable adjectives have been distinguished in the literature: relative (gradable) adjectives like *tall* and absolute (gradable) adjectives like *full* (Unger 1975; Kennedy 2007). Following one of Kennedy's tests, we say that a gradable adjective is absolute if the following entailment holds, and that it is relative otherwise:

(9) *X* is the ADJECTIVE one. $\rightsquigarrow$ *X* is ADJECTIVE (generally speaking)
a. Relative: My glass is the tall one. $\not\rightsquigarrow$ my glass is tall (generally speaking)
b. Absolute: My glass is the full one. $\rightsquigarrow$ my glass is full (generally speaking)

For example *tall* is relative because the inference is *not* systematic (9-a), but *full* is absolute because the inference *is* systematic (9-b). For the purpose of our study we chose to include only relative gradable adjectives. Our reason to do so was the following: intuitively, absolute adjectives denote an endpoint on a scale. For example, *full* denotes the maximum extent to which a recipient can be filled with substance.[4] A relative adjective like *tall*, on the other hand, does not select any context-invariant standard on a scale. Instead, it refers to a context-sensitive standard. In the way we constructed our stimuli, we always asked participants to imagine people or objects standing in the "middle range" of two extreme regions denoted by "very high" and "very low" along the corresponding dimension. For absolute gradable

for *healthy* if judgments for the sentence "this person is healthy" can be unclear and ambivalent. However, we excluded multi-dimensional adjectives from our sample as far as possible (see the discussion of evaluativity below), in order to rule out borderline cases arising from the relative weight of competing dimensions.

[4]The maximum standard to consider a glass "full" is still context-dependent: for instance, McNally (2011) discusses how wine need not reach the top of the glass for it to be considered full. A glass half-filled with wine can be called "full of wine" in some contexts, thereby threatening the generalization in (9). We don't think that undermines the point we are making here, however, as further tests can be used to corroborate the classification of "full" as absolute. However, the reader is invited to replace *full* by *empty* in our examples, as "empty" appears to show less context-sensitivity (in relation to glasses at least).

adjectives, we would have to adapt the descriptions. For example: if we were to ask participants to consider how they would describe glasses in the middle range between those that are "very filled" and those that are "filled very little", we would very likely fail to target the borderline region for what counts as "full". Intuitively, the borderline region for "full" is a region that is close enough to the maximum degree to which a recipient can be filled. Another reason we had for not including absolute adjectives in this experiment was that we also included antonyms, and here again, we can expect antonyms of absolute adjectives to not behave exactly like antonyms of relative adjectives (see Burnett (2016) for discussion). In summary, we did not include absolute adjectives in our experiment mostly to have a homogenous set of descriptions and adjectives to test.[5]

Criterion 2: We excluded absolute gradable adjectives from our list.

(iii) *Non-Evaluative, One-dimensional*

As Kennedy (2013) notes, all relative adjectives appear to be subjective. That is to say, they seem to systematically allow for faultless disagreement: substituting *tall* for ADJECTIVE in (10), as in (10-a), does not necessarily imply that either Mary or Sue is wrong, but doing the same with *prime* as in (10-b) *does* imply that either Mary or Sue is wrong. While Mary and Sue can truly diverge on what heights they consider to be tall, there is an objective standard for primeness and therefore one of them has to be wrong.

(10) X thinks Z is ADJECTIVE but Y does not.

a. Subjective: Mary thinks Paul is tall but Sue thinks he is not $\not\rightsquigarrow$ either Mary or Sue is wrong.

b. Not subjective: Mary thinks this number is prime but Sues thinks it is not $\rightsquigarrow$ either Mary or Sue is wrong.

In agreement with Kennedy's generalization, all the adjectives in our list satisfy the test of subjectivity.

However, some subjective adjectives are also *evaluative*, while others are not. We use the category *evaluative* in the sense of Kennedy (2013). That is, we call an adjective *evaluative* if substituting it for ADJECTIVE in (11) results in a non-deviant sentence.[6]

(11) Z finds that X is ADJECTIVE-er than Y.

[5] Some authors consider that the status of what we call borderline cases for absolute adjectives is due to a different phenomenon from the one at play with relative adjectives. For instance Kennedy (2007) claims that calling an almost-full glass "full" is a manifestation of *imprecision*, but that there is in fact a sense in which such a glass is uncontroversially *not* full. By contrast, for relative adjectives, there would be no non-arbitrary way to settle the question for borderline cases for they are a manifestation of *vagueness* proper.

[6] See Sæbø (2009) who first proposed tests of this sort. This sense of "evaluative", although related, is more specific than Rett (2007)'s sense, who calls an expression "evaluative if it makes reference to a degree that exceeds a contextually specified standard".

a. Evaluative: Mary finds that Paul is smarter than Joe.
b. Not evaluative: # Mary finds that Paul is taller than Joe.

As pointed by Kennedy (2013), evaluativity is a kind of subjectivity, as it correlates with faultless disagreement. However, evaluativity as diagnosed in (11) generally implies multi-dimensionality. In the case of "smart", evaluativity appears to correlate with the availability of different possible *respects* or *criteria* for building a scale on which one later establishes a threshold. Thus, one could entertain a relativistic approach where *X is* ADJECTIVE *and* NOT ADJECTIVE is interpreted as *X is* ADJECTIVE *according to some criterion and* NOT ADJECTIVE *according to some other criterion* (see Kamp and Partee (1995) for this observation, and Solt, this volume, for further remarks on that). To forestall any such explanation of our results,

Criterion 3: We excluded any evaluative adjective from our list.

In agreement with this criterion, the majority of the adjectives in our sample are associated with a unique salient dimension of comparison (such as "age" for "old", "height" for "tall", etc). One exception in our list is "rich", for which we mention "wealth", but for which we hint at a one-dimensional representation by talking of "degree of wealth".

(iv) *Individual-versus Stage-level*

The last criterion we used is based on the distinction between *individual*-level predicates and *stage*-level predicates. After Carlson (1977), we call an adjective *individual-level* if substituting it for ADJECTIVE in (12) results in a deviant sentence, and *stage-level* otherwise.

(12) There are two NOUN ADJECTIVE.
a. Individual-level: ? There are two men tall.
b. Stage-level: There are two men happy.

The properties attributed by individual-level predicates seem to be more inherent to their argument than those attributed by stage-level predicates. Once again, a relativistic approach could try to explain the acceptance of contradictory descriptions under a reading like *X is* ADJECTIVE *on some occasions and* NOT ADJECTIVE *on some other occasions*. To avoid that, we also relied on the following criterion:

Criterion 4: We excluded any stage-level adjective from our list.

Sentence types: Objet-oriented and human-oriented

It is important to stress that we used the four mentioned criteria fundamentally as a heuristics to build our stimuli. Our intention was not to test for those properties, nor to rigorously control for them, but to avoid confounds such as the ones we just discussed. Furthermore, the 8 adjectives that we used were distributed into two types of sentences: 4 adjectives were predicated of a subject denoting a human being (*tall*, *rich*, *heavy*, *old*) and 4 adjectives were predicated of a subject denoting an object or a property related to a human being (*loud*, *fast*, *large*, *wide*). Note that this division is not

inherent in the meaning of the adjectives ("heavy" for instance is applicable to both persons and objects), but we imposed it arbitrarily to create variety in the comparison classes. We therefore did not treat the object-oriented/human-oriented distinction as a controlled, independent variable, for we did not expect that it should impact our main prediction. We only found it useful to have a balanced set of examples. We presented half the participants with the four adjectives of the former type *first* and the four adjectives of the latter type *second*; the other half of the participants saw the reverse order. This parameter was crossed with the type of negation (syntactic vs. polar) and each participant responded to *syntactic* descriptions built with either human-oriented or object-oriented adjectives exclusively.

3.3 Participants

We recruited 148 participants online and anonymously via the Amazon Mechanical Turk platform. There a link would redirect them to the Ibex servers, on which the experiment was developed and hosted. Before going through the actual trials, participants first had to complete a pre-questionnaire consisting of seven simple questions. They were also presented with a post-questionnaire that was used for the study on the syntactic vs antonymic negation. Those forms are reported in Appendix 2. Accuracy on the pre-questionnaire and on controls was very good and no participant was excluded.[7] Though we didn't explicitly require our participants to be speakers of English, near-ceiling accuracy shows that all our participants had a good understanding of the English language, which was a precondition to evaluate the sentences that we used as our stimuli. Each participant was assigned one of four groups: two groups of participants judged either 4 human-oriented or 4 non-human-oriented syntactic descriptions *before* judging polar descriptions (again, treated as fillers in this paper), and the two other groups judged them *after* the series of polar descriptions.

3.4 Results

The barplot in Fig. 4 reports the average acceptance of each of the four types of descriptions exemplified in Fig. 3.[8] The plot in Fig. 5 shows how the participants behaved regarding the acceptance of "and"- and "neither"-descriptions based on the number of trials in which they gave similar answers: few participants ever rejected both descriptions (left-most bar) or ever accepted the "and"-description while rejecting the "neither"-description (second bar from the left). The very few participants

[7] In this case, exclusion of participants would lead the regression models that we ran to not converge, because of an insufficient variability on the controls.

[8] The results were descriptively similar in all conditions. A graph presenting the results for each condition is included in Appendix 3.

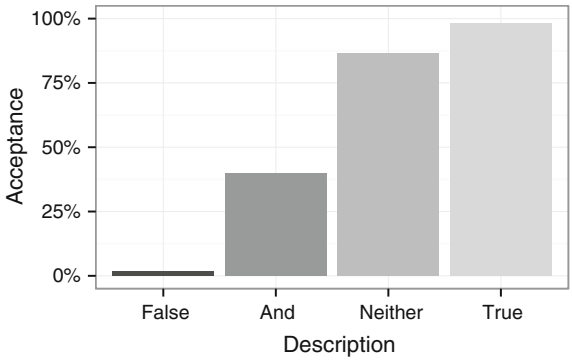

Fig. 4 Mean acceptance by description type

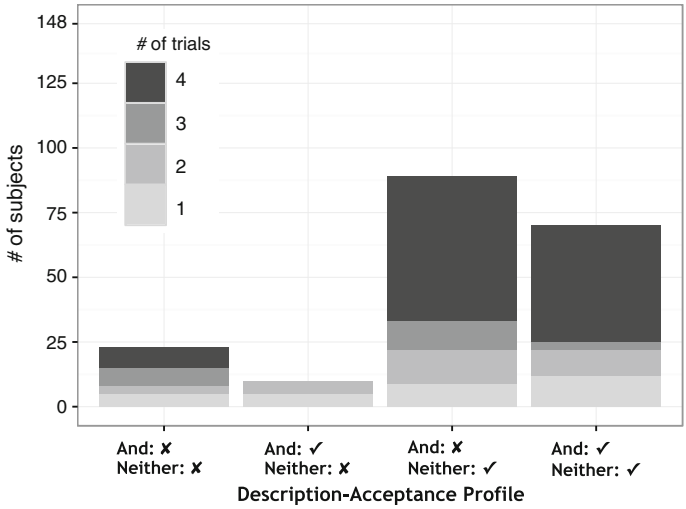

Fig. 5 Distribution of participants regarding the acceptance of "and"- and "neither"-descriptions based on the number of trials showing similar answers

who ever showed the latter behavior did it on only one or two trials, out of four. For most of the participants and most of the trials, the participants either accepted the "neither"-description while rejecting the "and"-description (third bar from the left), or they accepted both descriptions (right-most bar).

We ran mixed effect logistic regression models to analyze the data. We used the *glmer* function from the *lme*4 package (version 1.1–11) for R (version 3.1.2) with the optimizer "bobyqa" to compute the most complex models that would converge, following Barr et al. (2013). Several models with the same level of complexity would converge, but they all predicted the observation of a "Yes" response depending on Description (*Neither* vs. *And* vs. *Control True* vs. *Control False*) with random intercepts for Participant and Adjective. They differed in whether they also included

random slopes for Adjective, and whether they included random intercepts and random slopes for Block (*Before* vs. *After* judging descriptions with antonyms) or for Adjective Type (*Human* vs. *Non-Human* oriented).[9]

Throughout these models, we consistently found that the "and"-descriptions were accepted significantly more often than the control false descriptions, and significantly less often than the "neither"-descriptions. The models indicate that the "neither"-descriptions were accepted less often than the control true descriptions. We report the outputs of each model in Appendix 3.

3.5 Summary

The results confirm the previous observations from the experimental literature that we reported, according to which speakers do accept contradictory "neither"- and "and"-descriptions to describe borderline cases, given that those were accepted significantly more often than the control false descriptions. However, they provide us with a more nuanced picture of these judgments: first of all, it seems that neither of the two types of contradictory descriptions is generally as acceptable as plainly true (control) descriptions (a point that was left open in the existing experiments); and secondly, they show that "neither"-descriptions are significantly preferred to "and"-descriptions, eliciting a difference that was hinted at in the results of Alxatib and Pelletier (2011a) but not evidenced in those of Ripley (2011a). Note that the general preference for "neither"-descriptions over "and"-descriptions cannot be due to two different types of population, where one type of population (a majority) would accept only "neither"-descriptions and the other one (a minority) would accept only "and"-descriptions. Comparing the two bars corresponding to acceptance of "and"-descriptions in Fig. 5 (the second bar from the left and the right-most bar), it appears that for every participant, the acceptance of the "and"-description almost always co-occurred with the acceptance of the "neither"-description. By contrast, the "neither"-descriptions were often accepted alone (third bar from the left). As a consequence, on the one hand, these observations support the project of developing a system accounting for the acceptance of both kinds of descriptions but on the other hand, such a system should also account for the preference of "neither"-descriptions over "and"-descriptions.

4 Explaining the Asymmetry

Our two main predictions in undertaking this experiment were confirmed: first, we see that both "and"-descriptions and "neither"-descriptions are accepted significantly more than false control sentences using the same predicates, and secondly we see a

[9] See Appendix 3 for a list of the models.

marked preference for "neither"-descriptions over "and"-descriptions. This finding raises two main questions: the first is how we can explain the preference for "neither"-descriptions over "and"-descriptions, even as restricted to the class we considered in our study. The second is whether we can expect it to be robust across other kinds of adjectives than the ones we considered. In this section we focus on the first issue. We first discuss whether an explanation in terms of some domain-general bias is plausible. We propose instead a specific implementation of the strongest meaning hypothesis, based on the notions of strict versus tolerant meanings and the notion of local strengthening.

4.1 Omission Bias and Consistency Bias

One way in which one could be tempted to explain the preference for "neither"-sentences over "and"-sentences to describe a borderline case is as an instance of a domain-general bias toward omission rather than commission (see Spranca et al. 1991; Bonini et al. 1999). The explanation might go like this: a borderline case of a tall person is one for which participants feel uncertain whether to apply the predicate "tall" as opposed to "not tall". Participants may feel uncertain because, as postulated by Bonini et al. (1999), they may think that as a matter of fact only one of those descriptions is correct, but both descriptions compete on their mind. Intuitively, "neither"-descriptions can be taken to adequately express the participants' reluctance to ascribe either "tall" or "not tall". Instead of committing themselves to either the predicate or its negation, the participants express a preference for omitting both.

That explanation has a ring of truth, but it can't be quite adequate. One assumption that appears inadequate is that if participants felt that only one of the predicates "tall" and "not tall" ought to apply, then they should massively check "No" for descriptions like "x is tall and not tall". Unlike Alxatib and Pelletier (2011a), we did not leave participants with the option to say "Can't tell" in response to the sentences. But as our data indicate, however, about half the participants checked "Yes" to "and"-descriptions at least once (based on the "And: ✓" bars in Fig. 5). This confirms that an explanation in the style of Bonini et al. cannot be right: if we follow that explanation we can no longer explain why participants accept glutty descriptions to the extent that they do.

In light of our data, a different way of articulating the omission bias hypothesis may be as follows: participants do not feel that there is a fact of the matter dictating that only one description in terms of "tall" or "not tall" should be the correct one, but they may feel that once you commit yourself to one description, you should avoid using the other. In other words, participants may simply have a bias toward consistency. Although they feel that "tall" and "not tall" are equally applicable of a borderline case, they find more adequate to say that neither description applies, to avoid an inconsistency, rather than to say that both descriptions apply. In other words, the participants' behavior would simply reflect a preference for incompleteness over inconsistency.

That explanation sounds more convincing, but it still strikes us as ad hoc. Firstly, many participants accepted both kinds of contradictory descriptions on the same trial: in these cases, their bias toward consistency, which under this view is the reason why they accept the "neither"-description, would have to immediately fade in front of the "and"-description in order to allow for the acceptance of the latter. Secondly, it assumes that participants would interpret "tall" in a fixed way across both occurrences of "tall" and "not tall". What if participants had different interpretations in mind depending on the occurrence of "tall"? They may very well understand a sentence like "Betty is tall and not tall" to mean: "Betty is tall_1 and Betty is not tall_2", where "tall_1" picks a lower threshold for "tall" than "tall_2". If really participants think that no single way of drawing a line between "tall" and "not tall" people is correct, then this would be a very rational interpretation for them to use. Note that if this were the case, then "Betty is neither tall nor not tall" ought to mean "Betty is neither tall_2 nor not tall_1", which is logically equivalent to "Betty is tall_1 and Betty is not tall_2". Contextual variability easily avoids the inconsistency attached to either kind of description, but the preference for one kind of description over the other remains to be explained. If anything, following Grice's Maxim of Manner, one would expect to see a preference for the conjunctive description, since it is briefer and it appears morphologically simpler. In any event, appeal to a bias toward incompleteness over inconsistency loses its ground here.

In summary, we think that an explanation in terms of an omission bias can't be adequate, because it should predict that "and"-descriptions are not accepted at all. And an explanation in terms of a preference for incompleteness over inconsistency seems to us ad hoc and limited by assuming that participants would interpret "tall" rigidly, in a way that does not seem supported by the relative acceptance of sentences of the form "Betty is tall and not tall". Finally, a contextualist treatment can explain the consistency of borderline contradictions, but it does not straightforwardly explain the preference for "neither"- over "and"-descriptions.

4.2 Strict Meaning Versus Tolerant Meaning

In order to explain the findings of our study, we therefore turn to a distinct set of assumptions, and basically adopt the working assumption of Alxatib and Pelletier (2011a) and Cobreros et al. (2012) according to which vague adjectives are pragmatically ambiguous between two interpretations, a *tolerant* and a *strict* interpretation. This is similar to the idea of contextual variability, but putting more systematic constraints on the relation between two kinds of meaning.[10]

[10] Alxatib and Pelletier talk of sub- and super-interpretation. As pointed out by Cobreros et al. (2012), we can talk of sub- and super-interpretations, but provided we do not mistake the resulting logic of vague predicates for the subvaluationist and supervaluationist logics respectively, which are not truth-functional, unlike the strict-tolerant logic used in Cobreros et al. (2012). See Ripley (2013) and Alxatib et al. (2013) for discussion and comparison. The approach, while akin to the

We shall not review all the extant evidence for the distinction between strict and tolerant interpretations, but we only point out one of the earlier findings by Alxatib and Pelletier (2011a), which is that a significant proportion of their participants who checked True to the description of a borderline tall man as "tall and not tall" also checked False to the separate descriptions "tall" and "not tall".[11] The way that particular finding is explained by Alxatib and Pelletier (2011a) as well as Cobreros et al. (2012) is by appeal to the Strongest Meaning Hypothesis (SMH), namely the hypothesis that among two ambiguous meanings, the default should be to choose the strongest non-trivial meaning. In the case in question, the strongest non-trivial meaning that can be given to a sentence like "Betty is tall and not tall" is the conjunction of the tolerant meanings of "Betty is tall" and "Betty is not tall" respectively, whereas the strongest non-trivial meaning that can be given to the conjuncts separately is the strict meaning of "Betty is tall" and "Betty is not tall". For a speaker obeying the SMH, it is therefore consistent to accept "Betty is tall and not tall", while separately rejecting "Betty is tall" and "Betty is not tall", because the interpretation of "tall" and "not tall" switches from tolerant to strict from the conjunction to the conjuncts.

What the SMH encapsulates is a general bias toward the most informative meaning. We think the SMH can also be used to account for the preference of "neither"-descriptions over "and"-descriptions in our experiment. Our basic idea is simple: the idea is that when asked to decide whether a sentence containing a positive adjective or its negation is true or false, participants should have a bias toward selecting the strict meaning of the adjective and of its negation first. Only secondarily will they consider the tolerant meaning. To get the details right, however, we need to spell out some assumptions.

4.3 Local Strengthening

We adopt the three-valued presentation of strict and tolerant meanings given in Cobreros et al. (2015a). Given a propositional or first-order language, and a three-valued model, we call a sentence strictly true if it takes the value 1, and tolerantly true if it takes a value at least $\frac{1}{2}$ in that model, and we adopt the strong Kleene rules for the connectives. Given a vague predicate P, we represent the fact that a is a borderline case of P by assigning the value $\frac{1}{2}$ to Pa in that model. We now state our specific assumptions in order to explain the preference of "neither"-descriptions over "and"-descriptions.

(i) *Local Operators*

Our first assumption is that predicates are *locally* interpreted as tolerant or as strict at the subsentential rather than the sentential level. We basically posit strict and tolerant

contextualist strategy outlined at the end of the previous section, involves no mechanism of indexing. See Ripley (2011b) for more on this and the varieties of contextualism.

[11] See also Égré et al. (2013) for discussion, where a related phenomenon is discussed under the name "Hump Effect". The term "conjunction effect" now strikes us as more general and adequate.

Table 3 Truth-tables for the S and T operators

ϕ	$T\phi$	$S\phi$
0	0	0
$\frac{1}{2}$	1	0
1	1	1

operators S and T whose semantics is defined as in Table 3. The S and T operators correspond to Łukasiewicz's necessity and possibility operators in three-valued logic (see Malinowski 2007 for an overview). The S operator also corresponds to Bochvar's meta-assertion operator, sometimes written A, t or B (Bochvar 1937; Horn 1985; Beaver 2001; Spector 2012), which plays a role in the theory of presupposition projection (more on this in Sect. 4.4).

Local modulation of strength appears to us as a natural idea, given that strict and tolerant meanings seem to have adverbial reflects.[12]

(13) Betty is (somehow) tall and (somehow) not tall

(14) Betty is neither (clearly) tall nor (clearly) not tall

We do not say that *somehow* and *clearly* literally correspond to the strict and tolerant operators that we just introduced. Rather, what (13) and (14) show is that local modifications of the adjectives are a productive linguistic operation. In the rest of the discussion, we will assume the existence of two covert linguistic operators *strictly* and *tolerantly* that can appear in place of *somehow* and *clearly* and that respectively have the effects of T and S.

The algorithm must be such as to output the following interpretations for the above sentences[13]:

(15) $T(tall(a)) \wedge T(\neg tall(a))$

(16) $\neg(S(tall(a)) \vee S(\neg tall(a)) \equiv \neg S(tall(a)) \wedge \neg S(\neg tall(a))$

The assumption of local strengthening is a significant departure from most recent accounts of borderline contradictions, where the pragmatic meaning of a sentence is computed globally for the whole sentence (see Cobreros et al. 2012; Alxatib et al. 2013; Cobreros et al. 2015b). It is however suggested by (Kamp and Partee 1995, 156) in their discussion of the assertibility of literal contradictions and literal tautologies involving multi-dimensional predicates. They consider that in a sentence

[12] See for example Serchuk et al. (2011), who call "confusion hypothesis" the hypothesis of a systematic strengthening of vague adjectives by a covert "definitely" operator, following the terminology of Williams (2006) (based on a expression first due to P. Greenough).

[13] Since the algorithm we propose operates on high-level linguistic representations, it equally derives the expected interpretation regardless of the form of the logical translation of the *neither* descriptions.

such as "Bob is a man and not a man", each occurrence is interpreted differently, "as if it were modified by something like "in some respects"". We discuss the locality assumption as well as the quantification over respects in greater detail below.

(ii) *Predicate Negation*

To flesh out assumption (i), we supplement it with another one, which concerns the behavior of negation. We take it that the operators do not get embedded inside a negated predicate (as in *not tall*) except if specifically marked (as in *not* $\underline{tall}_{Focus}$ where the short break between *not* and *tall* might signal the presence of the *strictly* operator).

(iii) *Bottom-up Strengthening and Backtracking*

Our third assumption is that strongest meanings are computed incrementally in the course of building the syntactic representation of a sentence. The idea is that first, the leaves of a syntactic tree are given the strongest meanings. Then, given two subconstituents of a larger constituent, their meaning composition gets a check for nontriviality. We call a meaning trivial if it is necessarily empty or necessarily tautological. If it is nontrivial, the algorithm proceeds according to classical rules in order to deliver a semantic verdict relative to the model at hand. If a triviality is reached, then one needs to backtrack and reassign the leaves of the tree the next strongest meaning available in order to reiterate the algorithm, until the algorithm ends and gives a semantic verdict.

(iv) *Least Effort*

We note that Kamp and Partee envisage the acceptability of a sentence like "Bob is a man and not a man" as involving backtracking in a similar way: for them, the access to different respects is based on perceiving a contradiction from a uniform meaning, and on the need to abide by Gricean maxims. In order to explain the preference for "neither"-descriptions of borderline cases over "and"-descriptions, however, we need an additional assumption, which is that the simpler of two computational procedures should be generally preferred to the more complex. Or to put it differently, if a run of the algorithm involves backtracking, it will involve a more costly representation of meaning, and participants will be less likely to compute it. That is, if a triviality is reached in the course of building the meaning of a sentence, a participant can always be lazy and deliver a verdict according to the interpretation reached, instead of repairing to get a nontrivial meaning. Fundamentally, this means that backtracking is optional, an assumption which remains compatible with Gricean principles.

Illustration

Let us see how our algorithm works on the non-adverbial versions of sentences (13) and (14). The sentence (13) should first be interpreted as *Betty is strictly tall and strictly not tall* ($S(tall(a)) \wedge S(\neg tall(a))$) in virtue of (i), (ii) and (iii). That sentence is a contradiction, hence a trivial sentence. The next nontrivial interpretation we can obtain after backtracking is *Betty is tolerantly tall and tolerantly not tall* ($T(tall(a)) \wedge T(\neg tall(a))$) in virtue of (iii). Note that according to (iv),

given the model we assumed, participants will either be lazy and judge $S(tall(a)) \wedge S(\neg tall(a))$ to have value 0 relative to their model of the situation, or they will have worked out the meaning of the sentence to be $T(tall(a)) \wedge T(\neg tall(a))$, and they will give it the value 1. The case of sentence (14) on the other hand involves no backtracking at all. The meaning we get for the sentence is *Betty is neither strictly tall nor strictly not tall* ($\neg S(tall(a)) \wedge \neg S(\neg tall(a))$), which is nontrivial and gets the value 1 in the model. Importantly, this implementation naturally accounts for the quasi-absence of acceptance of the "and"-descriptions to the exclusion of the "neither"-descriptions (second bar from the left in Fig. 5). Participants who are willing to rescue an "and"-description from contradiction must do more effort than they have to do for a "neither"-description, while representing the strict meaning in both cases. We predict as unlikely, therefore, for participants to accept the former kind of description while rejecting the latter.[14]

4.4 Comparisons

4.4.1 Global Strengthening

In order to establish whether our algorithm is plausible, we first need to compare it with extant algorithms, and then to see whether it makes further adequate predictions. To the best of our knowledge the closest kin to our algorithm is sketched in remarks made by Alxatib and Pelletier (2011a) about the computation of sub- and super-interpretation. Although Alxatib and Pelletier do not outline a general algorithm, they make some suggestive remarks, for instance concerning the meaning of double negations of gradable adjectives. For example, a sentence like:

(17) Betty is not not tall.

Can be used to convey that Betty is not short, but also that she is not "tall *tall*", in other words that she is borderline tall. Our algorithm can derive that meaning, since we basically get *Betty is not strictly not tall* ($\neg S(\neg tall(a))$) as the strengthened meaning.[15] This means exactly that Betty is tolerantly tall. As already mentioned, a globalist algorithm like that in Cobreros et al. (2012) cannot derive that interpretation, since $\neg\neg tall(a)$ is necessarily equivalent to $tall(a)$ both under its strict and under

[14] 10 participants did accept the "and"-description while rejecting the "neither"-description at least once (5 participants did so on one trial, 5 participants did so on two trials). This could be noise produced by inattention on these trials, even though our participants were highly accurate. It is also conceivable, in principle, that on one or two trials those 10 participants exceptionally went through the whole procedure described above for the "and"-description but stuck with a classical, logically contradictory interpretation for the "neither"-description.

[15] Note that in (17), *strictly* would again appear under the first *not* and typically trigger a short break in the prosodic contour. Relatedly, when we paraphrased the meaning of (17), we used "not tall *tall*" to mean "not strictly tall". The repetition of the adjective seems to be another focus-related strategy to embed the *strictly* operator under negation (see (ii) above).

its tolerant interpretation. The same holds of Cobreros et al. (2015b)'s modified algorithm, which introduces no pragmatic difference between a sentence and its double negation. Our assumption that predicate negation and sentential negation should be treated differently plays a crucial role here.

Another problem that has been raised by Alxatib et al. (2013) for the original globalist account in Cobreros et al. (2012) concerns sentences such as:

(18) Betty is tall and Betty is not tall, or Betty is rich.

When interpreted strictly, the sentence means merely that Betty is strictly rich. But intuitively, it in fact says that Betty is borderline tall or clearly rich, which is stronger. Cobreros et al. (2015b) propose a more elaborate algorithm capable of deriving that meaning. Our algorithm also derives it straightforwardly: the incremental processing of (18) will treat the constituent subsentence "Betty is tall and Betty is not tall", but as illustrated earlier embedding the strict operator would yield a trivial constituent, so backtracking reinterprets to mean that Betty is tolerantly tall and tolerantly not tall, and "Betty is rich" will be interpreted strictly, and their union will be nontrivial.

We note finally that none of the other algorithms of pragmatic meaning we cited would predict a preference for "x is neither P nor not P" over "x is P and not P". Alxatib et al. (2013)'s algorithm based on fuzzy logic, for example, would predict (13) and (14) to be equally acceptable, like Cobreros et al. (2012) or Cobreros et al. (2015b). The difference lies mostly in the fact that pragmatic strengthening in those cases is done globally, and not by the insertion of local operators.

4.4.2 Local Accommodation for Presupposition

As pointed out to us by Benjamin Spector, the use of the local operators S and T bears a strong analogy with the use of accommodation operators in the theory of presupposition projection. Thus, Spector (2012) shows how local accommodation in a three-valued treatment of presupposition can be handled by means of the S operator. The operator returns the value 1 when a sentence is true (gets the value 1) and the value 0 when the sentence is false or undefined (gets the value 0 or #). Spector's observation is that one can find pairs of sentences that have exactly the structure of conjunctive versus disjunctive borderline contradictions, except that they involve presuppositional instead of (merely) vague expressions. His example is the following (Spector 2012, 2016):

(19) a. John stopped smoking and John did not stop smoking.
b. Neither John stopped smoking, nor did he not stop smoking.

Spector notes that (19-b) is acceptable in a situation in which John never smoked, in a way that (19-a) isn't.[16] Indeed, this appears to be what (19-b) means, namely

[16](19-a) may be acceptable if John is a borderline case of someone who stopped smoking. But set aside the vagueness of "stop" and "smoke" (assume they are fully crisp predicates).

that neither description is applicable to John, because he never smoked. The proper way to account for the acceptability of (19-b) is in terms of local accommodation, for which Spector's analysis is as follows:

(20) not (*S*(John stopped smoking) or *S* (John didn't stop smoking)).

When John never smoked, "John stopped smoking" is undefined, and by the semantics of *S* the whole sentence gets the value 1. Contrariwise, (19-a) stays false or undefined, however we apply the accommodation operator to its second conjunct, and irrespective of the value of the first conjunct. In other words, the sentence is subject to no pragmatic repair, unlike (19-b).

The analogy is indeed striking, but there remains an important difference between vagueness and presupposition, which is that conjunctive descriptions are acceptable with vague predicates in borderline cases to at least some extent, whereas they seem always unacceptable in cases of presupposition failure (see Spector 2012; Zehr 2014 for more on the comparison). In Spector's account of the presuppositional case, what matters is that (19-a) can never get the value 1, whereas (19-b) can. Interestingly, Spector (2012) mentions the potential usefulness of a dual operator such as *T* to handle vague sentences.[17] This is exactly what happens under our assumptions, since "Betty is tall and not tall" can get the value 1 as per the use of *T*. Regarding our main finding, however, the key part in our explanation lies in the supposition that the pragmatic process needed to give the sentence that value is more costly than it is for "Betty is neither tall nor not tall". So we acknowledge the analogy between Spector's pair and pairs consisting of conjunctive versus disjunctive borderline contradictions, but we think more is needed to account for the relative rather than absolute preference of one kind of sentence over the other in the vagueness case.

4.4.3 Quantification Over Standards

We deliberately set aside multidimensional adjectives from our sample of adjectives since "Betty is healthy and not healthy" is easily interpreted to mean that Betty is healthy in some respect (for instance blood pressure), and not healthy in some others (for instance cholesterol) (see Kamp and Partee 1995). Prima facie, a shift in respects of comparison is not what is operative with one-dimensional adjectives like "tall". However, the tolerant interpretation of a sentence like "Betty is tall and not tall" does convey that Betty is tall by some (acceptable) standard, and not tall by some (distinct, but equally acceptable) standard, relative to the same respect of comparison.

[17] Spector writes the operator in question *W*, for "weak truth", and writes the dual *B*. Note that we developed our account independently of Spector's, that is from the vantage of the strict-tolerant account of vagueness, without thinking about presupposition accommodation. Spector was motivated primarily by the phenomenon of local accommodation, and with no heed to the specific asymmetry between "neither" and "and"-sentences we discuss in the vagueness case.

And conversely, "Betty is neither tall nor not tall" appears to convey that not every standard makes Betty tall, and that not every standard makes Betty not tall.[18]

Galit Sassoon points out to us that Shamir (2013) found that one-dimensional adjectives "are not as good as multidimensional adjectives with modifiers over respects (such as e.g., "in some/most/every respect" or "except in some respects"), but they are not crushingly bad either, and their interpretation seems to build on quantification over standards of membership" (p.c.). That is, "tall in some respect" can be used to mean "tall by some standard", and "tall in every respect" "tall by every standard". Moreover, Sassoon (2013) found that in the case of multidimensional adjectives, "positive" adjectives (like "safe") are predominantly universally modified, but "negative" adjectives (like "dangerous") are predominantly existentially modified, even though either type admits both interpretations. Recently, Sassoon and Fadlon (2016) noticed that the "positive" one-dimensional adjectives that they used as fillers in their experiment show a similar preference for strong modifiers like "all" over weak modifiers like "some". For Sassoon, this fact coheres with our finding and account of the data.

We note that if "tall" predominantly means "by every standard, tall", then "not tall" could end up meaning "by every standard, not tall" or "tall, but not by every standard", depending on whether negation should take narrow scope or wide scope over "every". Either way, "tall and not tall" would be predicted to be anomalous under the universal interpretation of its first occurrence. On the other hand if "tall" can be existentially modified to mean "by some standard, tall", as we get by the insertion of the T operator, then "tall and not tall" becomes admissible again (provided the existential modifier takes wide scope over negation, as we postulate for T). Sassoon's observation raises a further question, however, which is whether multidimensional negative adjectives like "dangerous", whose predominant reading is existential over respects, might invert the pattern we observed in one-dimensional adjectives, that is, whether we might see an interaction between the predominant reading of the adjective (existential vs. universal over respects), and the relative acceptability of "and"-contradictions over "neither"-contradictions. For instance, would "this is dangerous and not dangerous" be relatively more accepted than "this is neither dangerous nor not dangerous" compared to "this is safe and not safe" relative to "this is neither safe nor not safe"? This is an interesting question to ask, which we leave for further investigation.

[18] And indeed, while the account of strict and tolerant we adopted here is cashed out in trivalent terms, the original account of Cobreros et al. (2012) defines the tolerant meaning of an adjective in terms of an existential quantification and the strict meaning in terms of a universal quantification, not directly over acceptable standards, but in ways that are intertranslatable with that approach.

5 Conclusions and Perspectives

We reported on two main findings in this paper. The first is a new confirmation of the fact that classical contradictions both of the "neither" type and of the "and" type are accepted and used by naive speakers to describe borderline cases. This confirms that both gaps and gluts are operational in the representation of borderline cases (see Égré et al. 2013 for a related point). The second and more interesting finding is that for a representative class of relative gradable adjectives, "gappy" descriptions of the form "neither P nor not P" are preferred to "glutty" descriptions of the form "P and not P". The sense in which "neither"-descriptions are "gappy" is that they express that a particular case falls in the underlap between two strict extensions, and the sense in which "and"-descriptions are "glutty" is that they express that the same case falls in the overlap between two tolerant extensions. In agreement with the earlier accounts of the pragmatic meaning of vague predicates presented in Alxatib and Pelletier (2011a), Cobreros et al. (2012), Égré et al. (2013), Cobreros et al. (2015b), we have argued that we can account for that preference if indeed there is a bias toward selecting the strict meaning first, but moreover, and more centrally regarding the implications of our the theory, if pragmatic strengthening is done locally rather than globally.

Several questions remain. One is whether we can obtain independent confirmation of our explanation of the preference for "neither"-descriptions over "and"-descriptions in terms of the latter involving more steps of computation. Our account has implications regarding the online processing of these sentences. An interesting avenue for future research would be to collect response times and eye-tracking data: we expect slower response times, and possibly longer fixation and more regression eye-movements when accepting the "and"-descriptions than when accepting the "neither"-descriptions, under the assumption that the backtracking steps that we posit to access a tolerant reading have the same effects as found in the interpretation of garden-path sentences (Frazier and Rayner 1982).[19]

Another issue concerns the generality of our finding for other adjectival types. Does our algorithm predict the same preference for "neither"-descriptions over "and"-descriptions for absolute adjectives like "empty"? Relatedly, what are the *facts* for absolute adjectives? To answer the first question, we need a representation of absolute adjectives in our model. Burnett (2014, 2016) argues that the strict denotation of "empty" coincides with the classical denotation of "empty" (that is, it should denote the zero degree of being filled on the relevant scale), though the tolerant denotation of "empty" can include recipients with a tiny bit of stuff in them. Symmetrically, Burnett takes the tolerant extension of "not empty" to be identical to the classical extension, but she assumes the strict extension to be a proper subset of the classical extension (hence "not empty", read strictly, does not mean "strictly speaking, not empty", but something like: "clearly not empty", as Burnett stresses). The strict/tolerant duality that we observe for relative adjectives is therefore preserved for absolute adjectives: the strict extension of "not empty" is the complement

[19]Thanks to G. Sassoon for making that suggestions.

of the tolerant extension of "empty" and conversely. As a result, and as for relative adjectives, the gap defined by the strict extensions of "empty" and "not empty" corresponds to the glut defined by the tolerant extensions of "empty" and "not empty". Using our algorithm, Burnett's account would therefore predict the same preference for a description of the form "neither empty nor not empty" over a description of the form "empty and not empty" when describing a glass with a tiny bit of liquid in it. Although both are pragmatically non trivial and true, the former should be preferred, by the principle of using strict meaning first.

Regarding the second question, presently we cannot rule out the possibility for "and"-descriptions to be preferred to "neither"-descriptions for absolute adjectives, contrary to the predictions just discussed. If that were the case, one ought to question the interpretation of the second member of the "neither"-descriptions. It seems to us that a description of the form "neither empty nor not empty" has a natural interpretation paraphrasable as "neither empty nor, *strictly speaking*, not empty". Using our *strictly* operator, this reading amounts to "neither *strictly* empty nor not *strictly* empty", which is a manifest contradiction. However, the present remarks are based entirely on conjectures, and we need to confront them to actual data in order to make progress. We are in the process of running a separate study on absolute adjectives in order to gain further insights about such examples. Meanwhile, we think our results on relative adjectives already give us a compelling argument in favor of some form of local pragmatic strengthening in the computation of the meaning of sentences involving vague vocabulary.

Acknowledgements We thank the Agence Nationale de la Recherche under program TrilLogMean ANR-14-CE30-0010-01 for support, as well as grants ANR-10-LABX-0087 IEC and ANR-10-IDEX-0001-02 PSL* for research carried out at the Departement of Cognitive Studies of ENS. PE thanks the Swedish Collegium for Advanced Study who hosted him as a visiting fellow during the time this paper was written (Feb-June 2016). JZ thanks the National Science Foundation grant BCS-1349009 attributed to Florian Schwarz, and Florian Schwarz himself, for their support. We are particularly grateful to Galit Sassoon, to Benjamin Spector, and to three anonymous referees for detailed comments and suggestions. We also thank Sam Alxatib, Elitzur Bar-Asher Siegal, Pablo Cobreros, Chris Cummins, Matti Eklund, Vladimir Goranko, Yael Greenberg, Yossi Grodzinsky, Nat Hansen, Hans Kamp, Louise McNally, Marie-Christine Meyer, Peter Pagin, Roumy Pancheva, David Ripley, Robert van Rooij, Hans Rott, Uli Sauerland, Stephanie Solt, Steven Verheyen, Åsa Wikforss for various discussions. We also thank the participants of the "Vagueness and Three-valued Logics" course given by PE at the LLCC at the University of Jerusalem in 2015, as well as audiences in Madrid, Stockholm, Uppsala, and at the ESSLLI 2016 summer school in Bolzano (Course on "Trivalent Logics and Natural Language Meaning"). We warmly thank the editors, Elena Castroviejo Miró, Louise McNally, and Galit Sassoon for their editorial work and feedback.

Appendix 1: Scenarios

Scenarios for Human-Oriented Adjectives

Rich

A survey on wealth has been conducted in your country. In the population there are people with a very high degree of wealth, and people with a very low degree of wealth. Then there are people who lie in the middle between these two areas.

Imagine that Sam is one of the people in the middle range. Comparing Sam to other people in the population, is it true to say the following?

Tall

A survey on heights has been conducted in your country. In the population there are people of a very high height, and people of a very low height. Then there are people who lie in the middle between these two areas.

Imagine that Sam is one of the people in the middle range. Comparing Sam to other people in the population, is it true to say the following?

Old

A survey on age has been conducted in your country. In the population there are people whose age is very high, and people whose age is very low. Then there are people who lie in the middle between these two areas.

Imagine that Sam is one of the people in the middle range. Comparing Sam to other people in the population, is it true to say the following?

Heavy

A survey on weight has been conducted in your country. In the population, there are people of a very high weight, and people of a very low weight. Then there are people who lie in the middle between these two areas.

Imagine that Sam is one of the people in the middle range. Comparing Sam to other people in the population, is it true to say the following?

Scenarios for Object-Oriented Adjectives

Fast

A survey on people's cars has been conducted in your country. In the population, there are people who own very high speed cars, and people who own very low speed cars. Then there are people who own cars that lie in the middle between these two areas.

Imagine that Sam is one of the people owning a car in the middle range. Comparing Sam's car to the cars of other people in the population, is it true to say the following?

Large

A survey on people's houses has been conducted in your country. In the population, there are people who own houses with a lot of space, and people who own houses with very little space. Then there are people who own houses that lie in the middle between these two areas.

Imagine that Sam is one of the people owning a house in the middle range. Comparing Sam's house to the houses of other people in the population, is it true to say the following?

Loud

A survey on people's voice has been conducted in your country. In the population, there are people whose voice has a very high intensity, and people whose voice has a very low intensity. Then there are people whose voice lie in the middle between these two areas.

Imagine that Sam is one of the people whose voice lie in the middle range. Comparing Sam's voice to the voices of other people in the population, is it true to say the following?

Wide

A survey on people's feet has been conducted in your country. In the population, there are people with a very high foot breadth, and people with a very low foot breadth. Then there are people whose foot breadth lie in the middle between these two areas.

Imagine that Sam is one of the people with a foot breadth in the middle range. Comparing Sam's feet to the feet of other people in the population, is it true to say the following?

Appendix 2: Questionnaires

Pre-questionnaire

Before proceeding to the actual experiment, please answer these simple questions.	
There are 7 days in a week	❍ Yes ❍ No
Barack Obama is the current President of the USA	❍ Yes ❍ No
Abraham Lincoln was born in 2003	❍ Yes ❍ No
Marilyn Monroe died in 1780	❍ Yes ❍ No
Nicolas Sarkozy was one of the Presidents of the United States	❍ Yes ❍ No
California is part of the USA	❍ Yes ❍ No

Post-questionnaire

Please answer these few questions about the expriment. We are interested in what you actually remember at this point, so please do not reload the previous pages. Thank you.

Was there a scenario describing a population of pregnant women? ❍ Yes ❍ No

Were the descriptions that you saw in the first half of the experiment of a different form from those in the second half? ❍ Yes ❍ No

[Can you give an example of the type of descriptions in the first half?

I

Can you give an example of the type of descriptions in the second half?

I]

The last two questions and their input fields would appear only if the participants reported a difference between the descriptions in the first and in the second halves of the experiment.

Appendix 3: Results Per Group and Regression Models

We ran several models, differing with respect to the complexity of their random effect structures. They all included a random intercept, but no random slope for Participant (N = 148). As explained in the Design subsection of Sect. 3, each participant was assigned to one of four groups, determined by two factors. Block Order indicates whether the participant responded *before* or *after* judging another set of descriptions with antonyms, and Adjective Type indicates whether the participant responded to descriptions directly referring to a human being as their subject or to an object or a property of a human being.

Table 4 Outputs for models with the formula $Response = yes \sim Description + (1|Participant) + (1 + Description|Adjective) + (1 + Description|Block)$

Neither	Estimate	Std. Error	z value	Pr(>|z|)
(Intercept)	3.5971	0.3419	10.521	<2e-16 ***
True	4.2426	1.6634	2.551	0.0108 *
And	−4.3489	0.4941	−8.801	<2e-16 ***
False	−10.5246	0.7412	−14.200	<2e-16 ***
True				
(Intercept)	7.839	1.334	5.875	4.22e-09 ***
Neither	−4.242	1.377	−3.082	0.00206 **
And	−8.591	1.688	−5.090	3.58e-07 ***
False	−14.767	1.429	−10.330	< 2e-16 ***
False				
(Intercept)	−6.9275	0.6254	−11.077	< 2e-16 ***
And	6.1759	0.8515	7.253	4.08e-13 ***
Neither	10.5247	0.7349	14.321	< 2e-16 ***
True	14.7673	1.6375	9.018	< 2e-16 ***

We included a random intercept for Adjective (N = 8) in all our models. We were also able to fit models including a random slope for Adjective. One set of such models added a random intercept plus a random slope for Block Order, and another set of such models added a random intercept plus a random slope for Adjective Type. For all these models, we ran three versions: one with *Neither* as baseline for Description, one with *True* as a baseline and one with *False* as a baseline. The outputs of the models for the former structure (with random intercepts and slopes for Block Order) are presented in Table 4, the outputs of the models for the latter structure (with random intercepts and slopes for Adjective Type) are presented in Table 5. By removing the random slope for Adjective, we were also able to fit models including a random intercept and slope both for Block Order and Adjective Type with *Neither* as a baseline. The contrast with *Control True* remains significant (Table 6).

To further investigate the source of the significant contrasts between *Neither* and *Control True*, and with regard to the apparently mixed descriptive results in Fig. 6 above, we ran additional models on subsetted data. We first considered four subsets: the responses of all the participants in the human-oriented adjective groups, those of all the participants in the non-human-oriented adjective groups, those of all the participants who responded in the first block and those of all the participants who responded in the second block. All these models had a random intercept for Participant, and a random intercept plus slope for Adjective, Block Order and Adjective Type. None of them yielded a significant difference between *Neither* and *True* ($0.2 < p < 0.25$). We then subsetted the data to each minimal group of participants.

Table 5 Outputs for models with the formula $Response = yes \sim Description + (1|Participant) + (1 + Description|Adjective) + (1 + Description|AdjectiveType)$

Neither	Estimate	Std. error	z value	Pr(>|z|)
(Intercept)	3.4863	0.3649	9.556	< 2e-16***
True	3.5304	1.1144	3.168	0.00154 **
And	−4.1911	0.4507	−9.298	< 2e-16 ***
False	−10.1704	0.6343	−16.035	< 2e-16 ***
True				
(Intercept)	7.0164	0.9039	7.763	8.33e-15 ***
Neither	−3.5301	1.0223	−3.453	0.000554 ***
And	−7.7211	1.3308	−5.802	6.57e-09 ***
False	−13.7003	1.2777	−10.723	< 2e-16 ***
False				
(Intercept)	−6.6840	0.5847	−11.43	<2e-16 ***
And	5.9791	0.5519	10.83	<2e-16 ***
Neither	10.1703	0.5871	17.32	<2e-16 ***
True	13.7010	1.3338	10.27	<2e-16 ***

Table 6 Output for one model with the formula $Response = yes \sim Description + (1|Participant) + (1 + |Adjective) + (1 + Description|Block) + (1 + Description|AdjectiveType)$

Neither	Estimate	Std. error	z value	Pr(>|z|)
(Intercept)	3.6215	0.4070	8.898	<2e-16 ***
True	4.2418	1.8012	2.355	0.0185 *
And	−4.3742	0.5169	−8.463	<2e-16 ***
False	−10.5672	0.7154	−14.771	<2e-16 ***

The models failed to converge for the responses from the first block for the non-human oriented descriptions, and for the responses from the second block for the human-oriented descriptions. The other two models indicated a significant contrast between *Neither* and *Control True* ($p < 0.01$).

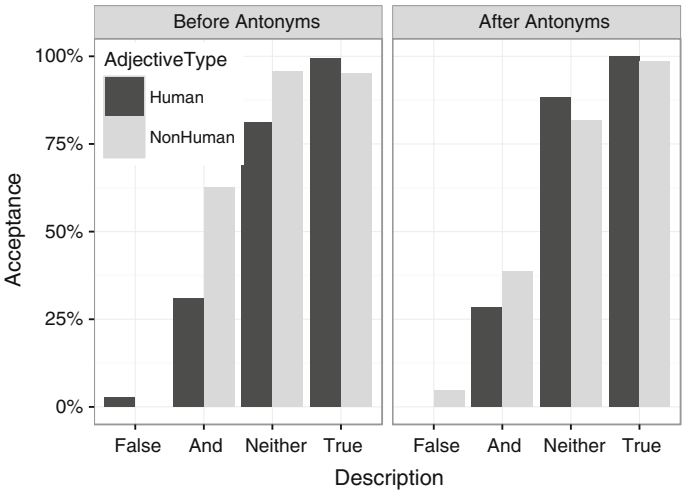

Fig. 6 Mean acceptance by description type and by group

References

Alxatib, S., & Pelletier, F. (2011a). The psychology of vagueness: Borderline cases and contradictions. *Mind & Language*, *26*(3), 287–326.

Alxatib, S., & Pelletier, J. (2011b). On the psychology of truth-gaps. In R. Nouwen, R. van Rooij, U. Sauerland, & H.-C. Schmitz (Eds.), *Vagueness in communication*. Lecture Notes in Computer Science (Vol. 6517, pp. 151–168). Berlin: Springer.

Alxatib, S., Pagin, P., & Sauerland, U. (2013). Acceptable contradictions: Pragmatics or semantics? A reply to Cobreros, et al. *Journal of Philosophical Logic*, *42*(4), 619–634.

Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, *68*(3), 255–278.

Beaver, D. I. (2001). *Presupposition and assertion in dynamic semantics*. Stanford, CA: CSLI publications.

Bochvar, D. A. (1937). On a three-valued calculus and its applications to the paradoxes of the classical extended functional calculus. *Mathematicheskii sbornik*, *4*(46), 287–308. English translation published in 1981 by M. Bergmann, *History and Philosophy of Logic*, *2*(1981), 87–112.

Bonini, N., Osherson, D., Viale, R., & Williamson, T. (1999). On the psychology of vague predicates. *Mind & Language*, *14*(4), 377–393.

Burnett, H. (2014). Penumbral connections in comparative constructions. *Journal of Applied Non-Classical Logics*, *24*(1–2), 35–60.

Burnett, H. (2016). *Gradability in natural language: Logical and grammatical foundations*. Oxford: Oxford University Press.

Carlson, G. N. (1977). *Reference to kinds in English*. Dissertation, UMass/Amherst.

Cobreros, P., Égré, P., Ripley, D., & van Rooij, R. (2012). Tolerant, classical, strict. *The Journal of Philosophical Logic, 41*(2), 347–385.

Cobreros, P., Égré, P., Ripley, D., & van Rooij, R. (2015a). Vagueness, truth and permissive consequence. In D. Achouriotti, H. Galinon, & J. Martinez (Eds.), *Unifying the philosophy of truth* (pp. 409–430). Springer.

Cobreros, P., Égré, P., Ripley, D., & van Rooij, R. (2015b). Pragmatic interpretations of vague expressions: Strongest meaning and nonmonotonic consequence. *Journal of Philosophical Logic*, *44*(4), 375–393.

Égré, P., De Gardelle, V., & Ripley, D. (2013). Vagueness and order effects in color categorization. *Journal of Logic, Language and Information*, *22*(4), 391–420.

Fine, K. (1975). Vagueness, truth, and logic. *Synthese*, *30*(3–4), 265–300.

Frazier, L., & Rayner, K. (1982). Making and correcting errors during sentence comprehension: Eye movements in the analysis of structurally ambiguous sentences. *Cognitive Psychology*, *14*(2), 178–210.

Horn, L. R. (1985). Metalinguistic negation and pragmatic ambiguity. *Language*, *61*(1), 121–174.

Kamp, H., & Partee, B. (1995). Prototype theory and compositionality. *Cognition*, *57*(2), 129–191.

Kennedy, C. (2007). Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, *30*(1), 1–45.

Kennedy, C. (2013). Two sources of subjectivity: Qualitative assessment and dimensional uncertainty. *Inquiry*, *56*(2–3), 258–277.

Malinowski, G. (2007). Many-valued logic and its philosophy. In *Handbook of the history of logic* (Vol. 8, pp. 13–94).

McNally, L. (2011). The relative role of property type and scale structure in explaining the behavior of gradable adjectives. In R. Nouwen, R. van Rooij, U. Sauerland, & H.-C. Schmitz (Eds.), *Vagueness in communication*. Lecture Notes in Computer Science (Vol. 6517, pp. 151–168). Berlin, Heidelberg: Springer.

Peirce, C. S. (1902). Vague. In J. M. Baldwin (Ed.), *Dictionary of philosophy and psychology* (p. 748). New York: Macmillan.

Rett, J. (2007). Antonymy and evaluativity. In *Semantics and Linguistic Theory* (Vol. 17, pp. 210–227).

Ripley, D. (2011a). Contradictions at the borders. In R. Nouwen, R. van Rooij, U. Sauerland, & H.-C. Schmitz (Eds.), *Vagueness in communication*. Lecture Notes in Computer Science (Vol. 6517, pp. 169–188). Berlin: Springer.

Ripley, D. (2011b). Inconstancy and inconsistency. In P. Cintula, C. G. Fermüller, L. Godo, & P. Hájek (Eds.), *Reasoning under vagueness: Logical, philosophical, and linguistic perspectives* (pp. 41–58). College Publications.

Ripley, D. (2013). Sorting out the Sorites. In K. Tanaka, F. Berto, & E. Mares (Eds.), *Paraconsistency: Logic and applications*, (pp. 329–348). Dordrecht: Springer.

Roche, L. (2012). *La négation des adjectifs*. Master's thesis, ENS/EHESS/Paris V.

Ruytenbeek, N. (2013). *An experimental approach of negated gradable adjectives*. Master's thesis, ENS/EHESS/Paris V.

Ruytenbeek, N., Verheyen, S., & Spector, B. (2017). Asymmetric inference towards the antonym: Experiments into the polarity and morphology of negated adjectives. *Glossa: A Journal of General Linguistics*, *2*(1), 92, 1–27.

Sæbø, K. J. (2009). Judgment ascriptions. *Linguistics and Philosophy*, *32*(4), 327–352.

Sassoon, G. W. (2013). A typology of multidimensional adjectives. *Journal of Semantics*, *30*(3), 335–380.

Sassoon, G. W., & Fadlon, J. (2016). *The role of dimensions in classification under predicates predicts their status in degree constructions*. Unpublished manuscript, Bar Ilan University.

Serchuk, P., Hargreaves, I., & Zach, R. (2011). Vagueness, logic and use: Four experimental studies on vagueness. *Mind & Language*, *26*(5), 540–573.

Shamir, A. (2013). *Disjunctivity*. Master's thesis, The Hebrew University of Jerusalem.

Solt, S., & Gotzner, N. (2010). *Expensive, not expensive, or cheap?* Paper presented at the 11th Szklarska Poreba Workshop.

Spector, B. (2012). Vagueness, (local) accommodation, presupposition and restrictors. Course Notes on Trivalent Semantics for Vagueness and Presupposition, Vienna.

Spector, B. (2016). Presupposition (and vagueness) projection at the propositional level. ESSLLI 2016 Course notes on Trivalents Logics and Natural Language Meaning, Day 2.

Spranca, M., Minsk, E., & Baron, J. (1991). Omission and commission in judgment and choice. *Journal of Experimental Social Psychology*, *27*(1), 76–105.
Unger, P. (1975). *Ignorance: A case for scepticism*. Oxford: Oxford University Press.
Williams, J. R. G. (2006). An argument for the many. *Proceedings of the Aristotelian Society*, *106*(1), 411–419.
Zehr, J. (2014). *Vagueness, presupposition and truth-value judgments*. Dissertation, ENS, Paris.