On the Psychology of Truth-Gaps*

Sam Alxatib[1] and Jeff Pelletier[2]

[1] Massachusetts Institute of Technology
[2] University of Alberta

Abstract. Bonini et al. [2] present psychological data that they take to support an 'epistemic' account of how vague predicates are used in natural language. We argue that their data more strongly supports a 'gap' theory of vagueness, and that their arguments against gap theories are flawed. Additionally, we present more experimental evidence that supports gap theories, and argue for a semantic/pragmatic alternative that unifies super- and subvaluationary approaches to vagueness.

1 Introduction

A fundamental rule in any conservative system of deduction is the rule of $\wedge$-Elimination. The rule, as is known, authorizes a proof of a proposition p from a premise in which p is conjoined with some other proposition q, including the case $p \wedge \neg p$, where p is conjoined with its negation. In this case, i.e. when the conjunction of interest is contradictory, $\wedge$-elimination provides the first of a series of steps that ultimately lead to the inference of q, for any arbitrary proposition q. In the logical literature, this is often referred to as the Principle of Explosion:

(1)	$p \wedge \neg p$	(Assumption)
(2)	p	(1, $\wedge$-Elimination)
(3)	$\neg p$	(1, $\wedge$-Elimination)
(4)	$p \vee q$	(2, $\vee$-Introduction)
(5)	q	(3, 4, Disjunctive Syllogism)

Proponents of dialetheism view the 'explosive' property of these deductive systems as a deficiency, arguing that logics ought instead to be formulated in a way that preserves contradictory statements without leading to arbitrary conclusions. One such formulation is Jaśkowski's DL [8], an axiomatic system that is adopted as a logic for vagueness by Hyde [7]. Hyde's reformulation provides a semantics for DL that relies on a system of precisifications[1]: a predicate P is

* We thank Paul Egré, James Hampton, David Ripley, Robert van Rooij, Phil Serchuk, the organizers and audience at the ESSLLI 2009 Vagueness and Communication worshop and the ENS Vagueness and Similarity workshop, and the two anonymous reviewers for their help and insightful input. Many further issues that have not been fully discussed in this paper are detailed in Alxatib and Pelletier (forthcoming) [1]. The research for this project was partially funded by F. J. Pelletier's NSERC grant #5525.

[1] Precisifications were first used in van Fraassen's work on presuppositions (cf. [5]).

R. Nouwen et al. (Eds.): ViC 2009, LNAI 6517, pp. 13–36, 2011.

associated with a set of classically-constructed 'sharpenings' (precisifications), each of which delineates a precise boundary between P's extension and its anti-extension. The semantics of Hyde's system is then set up so that a predicate P is considered to hold of an individual a iff a belongs to P's extension in at least one precisification. This creates a system that preserves what may in other logics be seen as inconsistencies, for it now becomes possible for P to simultaneously hold *and* not hold of a single individual, as would happen when the individual, say a, belongs to P's extension in one precisification, and to its anti-extension in another. In this case, $P(a)$ is said to fall into a truth-value 'glut'; since truth/falsity in this logic requires truth/falsity in at least one precisification, $P(a)$ would be true *and* false when a belongs to P in some precisifications but not in all.

Logics like Hyde's are called 'subvaluationary' logics. The truth-value gluts that are characteristic of these systems stand in contrast with truth-value *gaps*, which emerge in the *super*valuationary systems of Fine [4] and Kamp [9]. In these logics, which are also intended as logics of vagueness, truth/falsity is defined as *super*-truth/falsity, where super-truth is truth in *every* precisification, and super-falsity is falsity in every precisification. So, if an individual a belongs to P's extension in some but not all precisifications, the statement $P(a)$ will not be assigned any truth-value, for it is neither true in every way of sharpening P, nor false in every way of sharpening P.

Both families of logics are built on top of a system of precisifications, and in both logics the individual precisifications are respectful of classical predicate logic: every predicate within a precisification has an extension, and the complement of this extension is precisely the predicate's anti-extension. No individual is left behind[2]. It is only when truth is defined as super/sub-truth that borderline cases show non-classical properties, namely having two truth-values in subvaluations, and no truth-value in supervaluations. Both frameworks, however, share with classical logic the rule of $\wedge$-elimination: if $p \wedge q$ is true in some sharpening, then $p \wedge q$ is sub-true, and since the sharpenings are classically constructed, the sharpening in which $p \wedge q$ holds is a sharpening in which p holds and q holds. It follows, then, that there is a sharpening in which p is true, and there is sharpening in which q is true. This makes both p and q sub-true, and therefore true. The same can be said of a supervaluationary system: if $p \wedge q$ is super-true, then every sharpening is such that $p \wedge q$ holds in it, and because every sharpening is classical, every sharpening will be such that p holds in it and q holds in it. So p and q will be super-true, and therefore true.

Our goal in this paper is to show experimental evidence for a pattern that violates $\wedge$-elimination, and to show further that this pattern can be accounted for if both the sub- and the super-valuationary approaches are used together. In

[2] Actually, in many formulations of supervaluations (like Fine's for instance) there is mention of 'incomplete' precisifications. A precisification of a predicate is incomplete if its extension together with its anti-extension do *not* exhaust the domain of individuals in the model. But in these formulations, it is usually added that only complete precisifications are considered when evaluating whether or not a proposition holds, and this has the effect of making the system maximally faithful to classical logic.

the course of establishing our argument, we intend to show that the observations which we think reconcile the two approaches pose a considerable challenge to the epistemic hypothesis proposed in Bonini et al. (BOVW). In Sect. 2 we lay out the relevant theoretical foundations and provide a very brief description of the Sorites paradox, and of the solution claimed by supervaluationists, subvaluationists, and epistemicists. In Sect. 3 we describe BOVW's experiment and their epistemic interpretation of the data, and we argue against their criticism of gap-theories. In Sect. 4, we describe the experiment conducted for this study and show how the results pose problems for BOVW's view, and discuss in detail our interpretation of the data. Finally, in Sect. 5 we show what we think is evidence against $\wedge$-elimination, and propose a unification of sub- and super-valuations to account for it.

2 Background

Vagueness is most famously characterized as a logical problem in Eubulides's Sorites Paradox. In contemporary literature, the paradox is often formulated as an inductive proof of a false statement like (1c) from two unobjectionable premises like those in (1a) and (1b).

(1) a. A man standing 190 cm is tall.
b. A man who is just a millimeter shorter than a tall man is also tall.
c. A man standing 100 cm is tall.

Most of those who tackle the paradox concern themselves with the inductive step (1b). Fuzzy logicians like Machina [12], for example, observe that when one is afforded with an infinite number of truth-values, one can choose to assign the inductive step a truth-value just short of complete truth (hence its near-acceptability). To see how this resolves the sorites, consider a rewording of the inductive step as a conditional: if n is tall then $n-\delta$ is tall (for some small change δ). In many fuzzy logics, the truth value of a conditional is 1 iff the consequent is at least as true as the antecedent, and otherwise,

$$V(p \rightarrow q) = (1 - V(p)) + V(q) \tag{1}$$

Returning now to the sorites conditional, if we assign to its antecedent the truth-value n and to its consequent the value $n - \delta$, the conditional will turn out to be $(1 - \delta)$-true – just under completely true. The reason is that $(1 - n) + n - \delta$ will be $1 - \delta$, regardless of the value of n. If one were to apply modus ponens to this conditional together with the basic (completely true) premise (a) of the sorites, modus ponens will license a conclusion that is $1 - \delta$ true. But if we repeat the process, the next application of modus ponens will produce a conclusion that is slightly less true ($1 - 2\delta$ true), and as we advance down the height spectrum the conclusions will gradually become less true, so that by the time we get to 100 cm the truth of the conclusion will be much closer to falsity than to truth.

Like the fuzzy logician, the sub-/super-valuationist takes issue with the inductive step of the sorites. But on her account, the inductive step turns out false. Recall that sub-/super-valuationary semantics refer to precisifications, which are classical constructions. The conditions on truth, whether subtruth or supertruth, make the inductive step of the paradox false, for in *no* precisification is it true that small changes in degree go unnoticed; in every precisification, every predicate has a precisely defined extension, so in every precisification there is an n for which $P(n)$, but for which $\neg P(n-1)$. Since the inductive step is false in every precisification, it is sub-/super-false, and since it is sub-/super-false, it is false. Note, furthermore, that the inductive step is not true in *any* precisification, so it can never be true even under the subvaluationist's 'weaker' requirements.

The inductive step of the sorites is considered false in another view of vagueness: the epistemic view. Advocates of epistemicism, like Williamson [19] and Sorensen [16,17], dismiss the need for non-classical logics in the treatment of vagueness. They insist, instead, that *in reality* there is for each vague predicate a precise boundary that divides its extension from its anti-extension, but that the location of this boundary is unknown. In its defence of classicality, the view is similar to the supervaluationary approach, but it differs in that it claims a single, albeit unidentifiable, precise boundary for every vague predicate. This is the view that Bonini et al. claim to find experimental support for.

3 BOVW's Experiment

3.1 Method

BOVW administered in-class questionnaires (in Italian) to 652 students in Italian universities in two between-subject experimental conditions: True and False. We will follow BOVW and refer to the True group as the 'truth-judgers' and the False group as the 'falsity-judgers'. The two conditions had approximately the same number of students. The objective behind the questionnaires was to find, numerically, the boundaries that their subjects thought appropriate for attributing a vague predicate to a given entity/event. Participants were presented with scenario-question pairs such as the following example, using the vague predicate *tall* and the dimension of height. The difference between the conditions is highlighted by the italicized text. (English translations are taken from BOVW's paper).

A. Condition: TRUE

   When is it *true* to say that a man is 'tall'? Of course, the adjective 'tall' is true of very big men and false of very small men. We're interested in your view of the matter. Please indicate the smallest height that in your opinion makes it *true* to say that a man is 'tall'.

   It is *true* to say that a man is 'tall' if his height is greater than or equal to ___ centimeters.

B. Condition: FALSE

When is it *false* to say that a man is 'tall'? Of course, the adjective 'tall' is false of very small men and true of very big men. We're interested in your view of the matter. Please indicate the greatest height that in your opinion makes it *false* to say that a man is 'tall'

It is *false* to say that a man is 'tall' if his height is less than or equal to ___ centimeters.

Other items included the following: *mountain* (in terms of elevation), *old* (in terms of a person's age), *long* (in terms of a film's length), *inflation* (in terms of percentage), *far apart* (as between two cities, in kilometers), *tardy* (for an appointment, in minutes), *poor* (in terms of income), *dangerous* (cities, in terms of crimes per year), *expensive* (for 1300cc sedan cars), *high unemployment* (in percentage with respect to a country), and *populous* (for an Italian city, in population). In our study, we focus only on the adjective *tall*.

The data were collected through a series of studies, each differing (sometimes only slightly) in choice of predicate. BOVW also ran a set of studies in which the words 'true' and 'false' were removed from the query. In these questionnaires the instructions were modified as in (C) and (D), and were given to different participant groups than (A) and (B) above.

C. Condition: TRUE

When is a man tall? Of course, very big men are tall and very small men are not tall. We're interested in your view of the matter. Please indicate the smallest height that in your opinion makes a man tall.

A man is tall if his height is greater than or equal to ___ centimeters.

D. Condition: FALSE

When is a man not tall? Of course, very small men are not tall and very big men are tall. We're interested in your view of the matter. Please indicate the greatest height that in your opinion makes a man not tall.

A man is not tall if his height is less than or equal to ___ centimeters.

Following BOVW, we will refer to queries like (C) and (D) as the non-metalinguistic queries, in contrast with the metalinguistic queries seen in (A) and (B).

Theoretical Predictions

Before we show BOVW's results, we take a moment to outline what might be predicted by advocates of the three approaches we are focusing on: subvaluationism, supervaluationism, and the epistemic view.

To the subvaluationist, borderline cases are cases that fall in truth-gluts. So in a subvaluationary world, when one asks about the minimal value n that makes n tall (or makes it true to say that n is tall), one is asking for the n above which it is *subtrue* to say that n is tall. By definition, borderline cases qualify, because it is subtrue to say that a borderline case is tall. Similarly, when one asks about the m below which it is *false* to say that m is tall, one is asking about the m below which it is *subfalse* to say that m is tall, and again, this will include borderline cases. The subvaluationist therefore predicts n to be lower than m, that is, the responses of BOVW's truth-judgers should come out lower than those of the falsity judgers.

The supervaluationist predicts the opposite. Truth in this framework is super-truth, and falsity is super-falsity. So the lowest n that makes 'n is tall' true is the lowest n that makes it supertrue, i.e. makes n tall in every precisification. This will place n just above the borderline range because borderline cases will be excluded, for they are not tall in every way of making tall precise. The highest m that makes it false (i.e. superfalse) to say 'm is tall' will, for the same reasons, also exclude the borderline cases, and will land just below the borderline range. The prediction, then, is that the responses of the truth-judgers take a greater value than the responses of the falsity-judgers.

It is not entirely clear what the epistemicist might predict here, at least if he does not augment his view with one or more auxiliary assumptions. Indeed, BOVW add to their epistemic hypothesis an assumption that makes their predictions converge with those of the gap-theorist. We return to this after we show their findings.

3.2 Results

For almost every predicate they tested, BOVW find the average of the values provided by the truth-judgers to be significantly higher than that of the values provided by falsity-judgers. In the case of *tall*, for example, they find that the minimum height that makes a man tall – or makes it *true to say* that a man is tall – is higher than the maximum height that makes him not tall – or *false to say* that he is tall. The results from four of their six studies are shown in Table 1.[3]

While these findings contradict the predictions of glut-theories of vagueness, they seem to stand in support of gap-theories. Surprisingly, however, BOVW reject the gap account and instead promote the following epistemic hypothesis:

> Vagueness as Ignorance: S mentally represents vague predicates in the same way as other predicates with sharp true/false boundaries of whose location S is uncertain.

The reason that gaps appear, according to BOVW, is that speakers are in general more willing to commit errors of omission than commit errors of commission. In

[3] The predicate 'tall' was not used in their Study 3. Study 6, which did include 'tall', made explicit reference to the middle range, and is therefore excluded from the present discussion.

Table 1. Truth- and Falsity-judgments for '*n* is tall' (from BOVW)

	Study 1	Study 2	Study 4	Study 5
Truth-judgers	178.30 cm	179.55 cm	181.49 cm	170.28 cm
Falsity-judgers	167.22 cm	164.13 cm	160.48 cm	163.40 cm

other words, speakers would rather withhold the application of a predicate to an individual with an uncertain degree of membership than incorrectly ascribe the predicate to an individual of whom the predicate might not hold.[4] As a result, truth-judgers will provide the lowest value that they *confidently* think the predicate in question applies to, and falsity-judgers, likewise, will provide the greatest value that they *confidently* think the predicate does not apply to. The former value will of course turn out greater than the latter, and thus gaps emerge with *all* predicates, not just the ones that are usually seen to be vague.

The grounds on which BOVW reject the gap hypothesis, which otherwise seems a natural consequence of their empirical results, are predominantly theoretical. Their main points of criticism of gap theories are (1) that gap-theories do not offer an elegant account of higher-order vagueness, and (2) that, when examined in light of their data, gap theories lead to contradictory statements. We evaluate each of these grounds in turn.

Higher-Order Vagueness. Higher-order vagueness is the phenomenon that seems inevitable whenever one proposes that there is a 'gap' between the extension and the anti-extension of a predicate. For example, if one wishes to propose that, because there is no sharp cutoff line between the bald and the not-bald men, there must be a gap between the bald men and the not-bald men, filled by borderline-bald men, it seems impossible to then try to justify a sharp cutoff line between the bald men and the borderline-bald men either. Nor, on the other side of the gap, between the borderline-bald men and the not-bald men. So, there should be borderline cases of borderline cases: a 'second order vagueness'. But once a theorist starts down this path, it seems not possible to stop at all: there will be all levels of higher-order vagueness. Any rationale that could be given to stop at some particular high-order could have been used to not admit of the original first-order gap.

It does not seem like an easy task for the supervaluationist to provide an account of higher-order vagueness, since the framework, as we described it at least, allows three *sharp* possibilities: true, false, and neither. But Keefe [10] proposes the following maneuver: suppose borderlineness were to apply not only to the predicate itself, e.g. *tall*, but also to the admissibility of the way the predicate is made precise. When the admissibility of the precisifications is subject to borderlineness, one can imagine some individual a who is tall in every admissible precisification, but who is not tall in some precisification s of borderline

[4] Based on studies by Ritov and Baron [15] and Spranca et al. [18].

admissibility. In this case, we cannot say that a is super-tall, for he can only be super-tall if we ignore s, and we can only ignore s if it was *in*admissible. But we cannot say that a is borderline either (gappy that is), for that requires that a be not tall in some admissible precisification, and s is not quite admissible. This makes a a borderline-borderline case (2nd-order vagueness). If further conditions are imposed in higher metalanguage(s) on, say, the admissibility of admissibility, then finer gradations become more visible in the system, for that makes room for borderline-borderline-bordline cases, etc. We refer the reader to Keefe for more details.

BOVW's problem with this approach, and one of their reasons for rejecting gap theories, is that 'the mental representation of all these vague boundaries seems psychologically implausible' (p. 388). They add, furthermore, that if the ascent to higher orders of vagueness is stopped, the blur surrounding the gappy region will be replaced with a sharp line, and 'there is no introspective evidence for such a line' (also p. 388).

We officially suspend judgement on the issue of psychological plausibility. But we object to the way BOVW use introspection as a test of acceptability of a semantic theory. We note, as they do also, that there is no introspective evidence for the sharp but unknown divider that is presumed by their epistemic theory, a charge that BOVW address by saying that 'other semantic/conceptual principles have been plausibly ascribed to people who do not reliably acknowledge them' (pg. 387). So in considering the very same feature that their theory shares with an opposing theory, they happily cite this principle to defend theirs but will not consider it as a possible defense of the opposing theory. We think, therefore, that these 'psychological arguments' they use to favor their hypothesis and reject gap-theories are inconsistent.

The Absurdity of Denying Bivalence. BOVW begin their second argument against gap-theories by claiming that no difference was detected between the size of the metalinguistic gaps and the size of the non-metalinguistic gaps. Recall that BOVW used two survey styles, in one inquiring about the n for which it was *true* to say that predicate P holds of an individual (the metalinguistic questionnaire), and in the other inquiring about the n that makes an individual P (no mention of truth – the non-metalinguistic questionnaire). The comparisons for *tall* are shown in Table (2).[5]

If BOVW are right, the gap-theorist has to admit that the truth-conditions for 'n is tall' and ' "n is tall" is true' are the same, and similarly for 'n is not

[5] Indeed there seems to be no significant difference between the metalinguistic truth-judgements and the non-metalinguistic ones, but it is questionable whether the same holds of falsity-judgements; the average of the n for the metalinguistic falsity-judgers – taken as the average of Studies 1 and 2 – is 165.68 cm. For the non-metalinguistic studies, 4 and 5, the average comes to 161.94 cm. The difference between the two is 3.74 cm, which is almost 30% of what subjects, on average, claim to be the difference between '"x is tall" is true' and '"x is tall" is false'. It thus seems quite likely that there is a significant difference between metalanguage falsity and object language negation. We continue our reply, however, as if this difference was insignificant.

Table 2. Comparison of BOVW's metalinguistic and non-metalinguistic judgements

	Truth-judgements	Falsity-judgements
Metalinguistic (Studies 1, 2)	178.30 cm; 179.55 cm	167.22 cm; 164.13 cm
Non-metalinguistic (Studies 4, 5)	181.49 cm; 178.28 cm	160.48 cm; 163.40 cm

tall' and ' "n is tall" is false'. But BOVW argue against the viability of this position for gap theorists, as follows. Suppose height n is borderline tall. On a supervaluationary account, the statement 'n is tall' will have no truth value, that is, 'n is tall' is not true and 'n is tall' is not false. They give the following argument (pp. 388–389) to show that this cannot be correct (they wish the $\equiv$ to be read 'has the same truth conditions as'):

(1) 'n is tall' is not true (assuming n to be borderline)
(2) 'n is tall' is not false (assuming n to be borderline)
(3) n is tall $\equiv$ 'n is tall' is true (as shown by their experimental results)
(4) n is not tall $\equiv$ 'n is tall' is false (as shown by their experimental results)
(5) n is not tall $\equiv$ 'n is tall' is not true (from equivalence (3))
(6) n is not not tall $\equiv$ 'n is tall' is not false (from equivalence (4))
(7) n is tall $\equiv$ 'n is tall' is not false (double-negation in (6))
(8) n is tall (from assumption (2) and equivalence (7))
(9) n is not tall (from assumption (1) and equivalence (5))
(10) n is tall and n is not tall (conjunction of (8) and (9))

Since (10) is contradictory, and furthermore goes against the anti-glut findings of BOVW's experiments, the assumptions (1) and (2) must be revised. But these assumptions are the very ones that define the supervaluation position! So, unless there has been a mistake in the reasoning that got us from these two assumptions and the experimental results, it appears that supervaluation theory has been refuted.

We think that a supervaluationist could legitimately complain about the inferences involving negation in BOVW's proof. Before we discuss this, we point out that the proof need not be explained in full in order to understand how the alleged absurdity arises; one need only look at (4) and (5) to see the problem: (4) and (5) have the same proposition to the left of the '$\equiv$' symbol, but they each describe a different state of affairs on the right side of '$\equiv$'. In (4), 'n is not tall' is claimed to have the truth-conditions that make 'n is tall' false, but in (5), 'n is not tall' is claimed to have the conditions that make 'n is tall' not true. This is trouble for the gap-theorist because in her theory the conditions that make 'n is tall' false are different from those that make it not true; 'n is tall' is false whenever it is superfalse, but it is not true whenever it is either false or neither true nor false. The two scenarios cannot *both* be said to have the same truth-conditions as 'n is not tall', precisely because they describe different truth-conditions. If BOVW can show that the gap-theorist is forced to accept (4) and (5), their argument succeeds.

But the gap-theorist is not forced to accept (4) and (5) in the way intended by BOVW. (5) is derived from the equivalence in (3), which says that whenever n is tall, 'n is tall' is true. From this, it follows that whenever n is not tall, 'n is tall' is not true. (5), then, is to be understood as saying that whenever n is *anything but tall*, the sentence 'n is tall' is not true. Now if we turn our attention to (4), it simply says that, based on empirical evidence, the gap-theorist ought to say that the sentence 'n is tall' is false whenever n is not tall. In order for their argument to be convincing, BOVW must force the gap-theorist to say that 'not tall' in this context also means *anything but tall*, just like it does in (5). In other words, BOVW seem to be saying that, in order for the gap theorist to make her theory match the empirical findings, she must say that 'n is tall' is false iff n is *anything but tall*. But this is something that BOVW cannot do; the gap-theorist can respond by saying that 'n is not tall' in (4) means 'n is super-not-tall'. If the negation in the left-side of (4) is assigned a strong interpretation, the problem for the gap-theorist described in the previous paragraph disappears.[6]

Essentially, the gap-theorist's escape is to say that negation can have two interpretations: 'n is tall' is false whenever n is *strong*-not tall ('choice' negation), and 'n is tall' is not true whenever n is *weak*-not tall ('exclusion' negation). Ultimately, we will favor an account where negation is treated unambiguously in the semantics, but where its different interpretations arise from pragmatic principles (we refer the reader to Horn [6] and Levinson [11] for discussions of the inferences involving negation). For now, however, we use truth-tables (Table 3) merely to illustrate the difference between the strong/choice and the weak/exclusion interpretations of negation.

Table 3. Strong/Choice negation ($\sim$) and Weak/Exclusion negation ($\neg$)

φ	$\sim\varphi$	$\neg\varphi$
T	F	F
G	G	T
F	T	T

It can now be seen that, with the distinction between negations in place, the conclusion in (10) loses its contradictory reading; (10) becomes the proposition that (the borderline case) n is tall, in the sense of 'not *un*tall', as it were, and not tall, in the sense of weakly-not tall, or not definitely tall. In other words, the conditions under which (10) holds are the very conditions that make *tall* and *not tall* subtrue. The reader is invited to verify this claim.

[6] Further discussion against this and related 'logic of the argument' is given in more detail and against a wider group of similar arguments, in Pelletier and Stainton [13].

BOVW's concern, then, is that by virtue of its sub-tallness and sub-not-tallness, n can be said to be tall and not tall, which is contradictory. In Sect. 5, we reveal some empirical evidence not only that this 'contradictory' conclusion is often judged true, but that its conjuncts are often considered false at the same time. Before we get to that, however, we describe the experiment and discuss the findings that we think are problematic for the epistemicist.

4 Experiment

Participants. 76 undergraduates from Simon Fraser University participated in this study. 59 participants classified themselves as fluent English speakers, 10 as advanced, and 5 as intermediate (leaving 2 participants, who left the question unanswered).

Method. Participants were presented with an image of five suspects in a police line-up (Fig. 1). The suspects were shown with the following heights in pseudo-randomized order: 5′4″, 5′11″, 6′6″, 5′7″, and 6′2″.[7]. The suspects were labeled with numbers on their faces, and were referred to by these numbers in the experimental material.

Participants were given a paper and pen questionnaire (on a separate page from the image) with five sets of four statements (one for each suspect). Each statement had three labeled checkboxes to the right. An example is included in (2) for suspect #1.

(2)	#1 is tall	True ❏	False ❏	Can't Tell ❏
	#1 is not tall	True ❏	False ❏	Can't Tell ❏
	#1 is tall and not tall	True ❏	False ❏	Can't Tell ❏
	#1 is neither tall nor not tall	True ❏	False ❏	Can't Tell ❏

Before the survey was handed out, the participants were given the following instructions:

- You will be asked to describe the heights of the five suspects in the line-up shown below.
- Please use the height standards of adult males in present-day North America.
- This is not a test, and there are no correct answers. Upon reading the questions, simply check the first answer that pops in your head and seems to describe the situation as you see it.

In order to minimize the effect of order on the subjects' responses, each sheet was printed with the questions randomly ordered. This was done in every copy of the survey, so no two copies had the same order of questions.

[7] Both the metric measurement system and the imperial system are in common usage in western Canada.

Fig. 1. Suspects of Different Heights in Police Lineup

Results and Discussion. Our reply to BOVW draws particularly on the responses to the first two statements. Later, in Sect. (5.1), we consider the other two sentences, in the course of presenting our own position. In Fig. 2, the percentages for *true* responses to *X is tall* are shown to increase with height, starting with 1.3% at 5′4″, reaching the median value of 46.1% at 5′11″, and peaking at 98.7% at 6′6″.

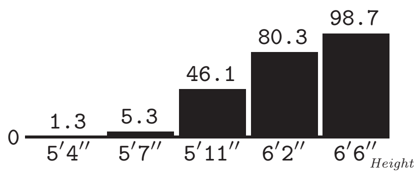

Fig. 2. % of 'True' responses to 'X is tall'

Conversely, the percentage of *false* responses, seen in Fig. 3, begins with a ceiling of 98.7% at 5′4″ and drops to 1.3% at 6′6″, passing the median at 5′11″ with a value of 44.7%.

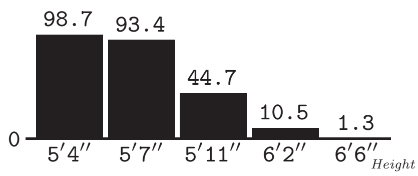

Fig. 3. % of 'False' responses to 'X is tall'

Figure 4 shows the percentage of *true* responses to *X is not tall*, which also reaches the median at 5′11″, this time at 25.0%, and peaks at 5′4″ at 94.7% and drops to 0.0% at 6′6″.

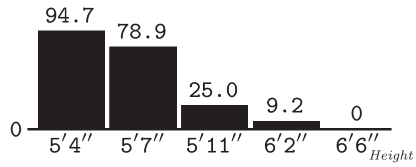

Fig. 4. % of 'True' responses to 'X is not tall'

The percentage of *false* responses to *X is not tall* is shown in Fig. 5: 3.9% at 5′4″, a median of 67.1% at 5′11″, and a maximum of 100.0% at 6′6″.

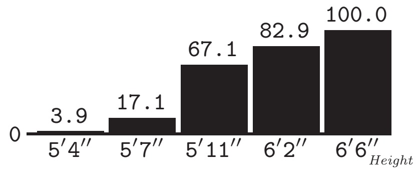

Fig. 5. % of 'False' responses to 'X is not tall'

It is the difference between these sets of answers that is problematic for the BOVW account. The numbers indicate a significant preference for rejecting a proposition over accepting its negation.[8] In classical logic, the statement '*a* is tall' is true just in case its negation, '*a* is not tall', is not true, and vice versa. But in a gap theory like supervaluations, the statement '*a* is tall' is true if it is supertrue, and otherwise it is not true. The prediction, then, is that if *a* is borderline, the statement '*a* is tall' is judged false more frequently than its negation '*a* is not tall' is judged true, the reason being that the latter statement only holds if it is supertrue, which would not be the case if *a* was borderline. Similarly, a gap theory would predict more false responses to '*a* is not tall' than true responses to '*a* is tall'.

Here we see an immediate objection: falsity in supervaluations is *super*falsity, and this disqualifies the tallness of borderline individuals from being false. A supervaluationist should not expect a preference for 'False' responses any more than a preference for 'True' responses when it comes to a borderline case. Strictly speaking, this objection is accurate. But of course, this would not be an issue if the checkboxes in our questionnaires were instead labeled 'True', '*Not true*' (instead of 'False'), and 'Can't tell'. For 'Not true' would surely include the borderline range for the gap theorist. But now suppose that a participant was in disagreement with a statement, and the only three options (as in our questionnaire) were 'True', 'False', and 'Can't tell'. We find it quite reasonable to expect the participant in this case to check 'False', since among the available answers, 'False' is the only plausible substitute for 'Not true'. We feel, therefore, that it

[8] According to a χ^2 test for independence, the chance of the difference (between denial and assertion) in the case of #2 being drawn from the same distribution is less than 5%: $\chi^2(2) = 8.22; p < 0.05$.

is legitimate to interpret 'False' as a sign of rejection in our set-up, but we certainly allow that this unsupported claim needs to be bolstered by experimental evidence.

The data, as shown in Figs. 2–5, confirms the preference for 'False' responses. For suspect #2 ($5'11''$), our borderline poster-child, 46.1% thought it was *true* that he was tall, while 67.1% thought it was *false* that he was not tall. Similary, 25.0% thought it was true that he was not tall, whereas 44.7% thought it was false that he was tall. Both comparisons show that a significantly bigger sample of participants rejected the statement '#2 is tall' (or 'not tall') when compared to the sample of those who accepted the classical negation of each statement.

BOVW might claim that this could as easily be taken as support for their epistemic hypothesis. Recall that BOVW assume that errors of commission are considered by their participants to be graver than errors of omission. Thus the subjects prefer to withhold judgement regarding uncertain cases than incorrectly attribute the predicate to them. If our participants (as we claim) would rather reject a statement (by judging it false) than accept its negation (by judging the negation true), can we not interpret this preference also as a way of favoring errors of omission over errors of commission? If this interpretation of the data is available, then the evidence that we find supportive of gap theories can also be taken to support the epistemic hypothesis (together with BOVW's auxiliary assumption regarding error preferences). In response to this concern, we point out that our subjects were also given the option of checking 'Can't tell', but very few people chose to answer that way: for the statement 'x is tall', where x is $5'11''$, there were 44.7% false responses, and 9.2% 'Can't tell' responses; for 'x is not tall', at the same height, there were 67.1% false responses, and 7.9% 'Can't tell's.

Of course, one may also object that it is possible for the participant to have taken 'Can't tell' as meaning something like 'I give up', thus accounting for the low rate of 'Can't tell' responses (because we cannot expect our subjects to comfortably choose this way of answering). In this picture, the fact that there is a preference for falsity-judgement over truth-judgement may after all be due to a preference of omission errors over commission errors, and so this part of our argument against BOVW is not convincing. Evidence against this interpretation is available elsewhere in our data, however. For, if we maintain vagueness-as-ignorance and combine it with this error-preference pattern, we should expect these same falsity-judgers (who are choosing to answer safely, as it were) to also prefer answering 'False' for the apparently contradictory statement '#2 is tall and not tall'. After all, the epistemic theory is classical, so it should predict virtually *no* 'True' responses to this statement. But as we will show below in Sect. 5, subjects seem happy to claim that this statement is true.

In a last-ditch attempt to save epistemicism, such a theorist may say that our last considerations cannot be taken as a counterargument to the vagueness-as-ignorance hypothesis because, the theorist might say, speakers need not be *aware* of their ignorance. This reply is not relevant here. What is relevant is that if errors of omission are indeed preferred to errors of commission, which is an assumption that the epistemicist requires, then we would expect a much larger number of 'Can't tell's, since this is the least committing answer with regards to borderline (or uncertain) cases.

We now return to the use of negation in this experiment and its role in the semantic/pragmatic account that we favor. Earlier we argued that BOVW were mistaken in assuming that only one type of negation could be understood in statements like 'a is not tall'. This assumption led them to conclude that ' "a is tall" is false' held under the same conditions as ' "a is tall" is not true', since both metalinguistic statements were 'equivalent' to 'a is not tall'. In response, we suggested that two interpretations are available for 'a is not tall': one in which the negation is identified with choice/strong negation (in which case 'a is not tall' holds if it 'super-holds'), and another in which the negation is identified with exclusion/weak negation (in which case the statement holds just in case 'a is tall' does *not* super-hold)[9]. A natural question that one can ask at this point is: which of these two types do we think arises when we present our participants with the statement 'X is not tall'? Surely, if the negation was interpreted as weak negation, then there should not be a significant difference between accepting the statement '#2 is not tall' and rejecting the statement '#2 is tall', since '#2 is not tall' (where 'not' is weak) would hold in the same set of circumstances that makes '#2 is tall' not hold. But since we do find a significant preference to deny the former, it would seem that the negation is interpreted as strong, and we should explain why.

We think that pragmatic factors contribute to the emergence of what resembles a strong/choice interpretation for 'not'. We begin our explanation of the pragmatic effects by inviting the reader to consider the following scenario: suppose John and Mary have a single friend named Lucy. Lucy is looking for a date, and John and Mary suggest that she meet their friend Bill. Suppose further that Bill is of average height. Now Lucy asks their friends about Bill's looks, and in response, Mary provides a few answers, one of which being 'he's not tall'. Here we find it felicitous of John to object to the way Mary described Bill's physical stature, and say in response: 'Well, he's not *not tall.* He's average.' The felicity of this interaction suggests that two different logical interpretations of

[9] We wish to note here that there is also room for interpreting negation as intuitionistic negation. The intuitionistic negation of p, $-p$, is true iff p is *false*, and is false when p is true or, on a gap-theoretic interpretation, when p is neither true nor false. In singly-negated statements, intuitionistic negation converges with choice/strong negation, since both assign the value True to $-p$ whenever p is false. However, intuitionistic negation differs from choice/strong negation in doubly-negated statements: $--p$ is true whenever $-p$ is false, and $-p$ is false iff p is true or truth-valueless. At first glance, this seems desirable, since we can derive the 'not untall' reading of 'not not tall' using only one definition of negation, and also derive the strong interpretation of negation in singly-negated expressions. But the consequence of having only intuitionistic negation in the language is that singly-negated expressions can *never* be given a weak interpretation. In Sect. 5.2 we show an example in which a weak interpretation of single negation is required (see Footnote 17).

'not' are involved, for otherwise John's comment would merely be equivalent to 'he's tall'.

If we assume that Bill is of average height, and if we are considering a gap-theoretic system, then the statement 'He is tall' will be neither true nor false. So, when Mary says 'He's not tall', if she is to be speaking truthfully[10] she must be intending a weak/exclusion interpretation of 'not', for that is the only negation that will convert an 'other-valued' statement into a truth. When John tries to correct Mary, or correct the impression left by Mary's statement, he takes what Mary said, with the truth value thus computed, and negates that. In this case, John is either negating a 'true' (if Mary was using weak/exclusion negation) or negating a 'other' (if Mary was using strong/choice negation). Note in these cases that if Mary were using weak/exclusion negation and were understood to be using weak/exclusion negation, then no matter what negation John is using to negate that, what he says is false, because all of the negations would take Mary's true into a false. We presume this is not right, since John is imagined to be speaking truthfully. From this it follows that, even if Mary were speaking truthfully (by using weak/exclusion negation), she could not have been understood that way. So it seems that John is taking Mary to be using strong/choice negation and he is denying the 'other' value to Mary's claim. This would mean that John was using weak/exclusion negation, since that is the only negation that maps an 'other' value to truth.[11]

So it seems that 'not' can be interpreted in a way akin to strong/choice negation in some cases, and to weak/exclusion negation in others. In order to provide a complete account of how negation is used, particularly with vague predicates, one must offer a description of the situations in which the strong/choice interpretation arises, and those in which the weak/exclusion interpretation arises. Here we follow Levinson (among others) and invoke familiar pragmatic principles: when we, as experimenters, present a group of participants wih questions or statements that contain negated (or even unnegated) vague expressions, we feel it reasonable to assume that these expressions are being interpreted by the participants with sufficient observance of the Gricean maxims, in particular the maxim of quantity. If it is also assumed by our participants that we intend for this principle to be observed by them, then we would expect that by '(not) tall' the participants will understand that we want them to pick up on the most informative reading possible, which to the participant must correspond to that definition of '(not) tall' which s/he thinks all (or most) people would agree upon and, also, that s/he assumes that we, the experimenters, think all (or most) people would agree upon (assuming, of course, a fixed context of use, comparison class, etc.). The closest match to this description is the super-interpretation, i.e.

[10] 'Speaking truthfully' here is to be understood as making a *semantically* true statement. The statement might not accord with various Gricean restrictions and therefore might not be a *pragmatically* felicitous statement.

[11] Here it is also possible to understand John as using his own negation twice, in which case it may be that John is in fact using intuitionistic negation. This possibility was brought up and discussed briefly in Footnote 9.

that 'is (not) tall' is read as 'is *super*-(not)-tall'. So, when the question arises as to whether a person standing 5′11″ is tall (or not tall) the addressee – who may reasonably be expected to comply with the Gricean principles – is very likely to say 'False'. In the next section we present more experimental findings (Sect. 5.1) and offer a more detailed theoretical account in which the maxims of quality and manner are involved (5.2).

5 Contradictions and Borderline Cases: Gaps vs. Gluts

In this section we turn to statements in our questionnaire that until now we have ignored: 'x is tall and not tall' and 'x is neither tall nor not tall'. The relevant data are by no means indicative of a knock-down argument in favor of any particular theory, but the implications they carry can be of great importance for the gap theorist as well as the glut theorist, and we intend to use them to further clarify and extend our account of the data we have already presented.

5.1 Data

Figures 6 and 8 show that the numbers of *true* responses to each of these statements, which we will call *both* and *neither*, increased when the suspect's height was closer to average, peaking at 44.7% and 53.9%, respectively, for the 5′11″ suspect.[12] The number of *false* responses followed a complementary pattern, *decreasing* as the heights approached 5′11″ and reaching a minimum of 40.8% and 42.1% at that midpoint, as shown in Figs. 7 and 9. Note that there are more subjects who say *true* to '#2 is tall and not tall' than say *false* to it. Note also that more subjects say *true* to '#2 is neither tall nor not tall' than say it is *false.*

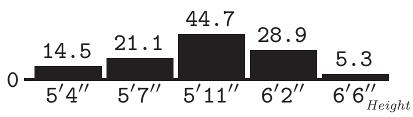

Fig. 6. % of 'True' responses to 'X is tall and not tall'

Particularly interesting, however, is how the two statements, *both* and *neither*, correspond with one another. The responses for the two questions are cross-tabulated in Table 4. Note that the response types are subscripted with the

[12] One may question the reliability of 'True' responses to the contradictory statement here. For example, it may well be that (relative) abundance of 'True' responses to *both* is due to a simple yes-bias. Regarding this concern, however, we find it unlikely for a yes-bias to increase the number of 'True' responses to the *both* statement and not to its individual conjuncts. This of course deserves further investigation. Experimentally, one could compare the frequency of truth-judgements to a statement like 'tall and not tall' to a stronger one like 'definitely tall and definitely not tall'. If we find significantly fewer truth-judgements to the latter, our findings will no doubt be more informative. We thank James Hampton and Lawrence Goldstein for bringing up this issue.

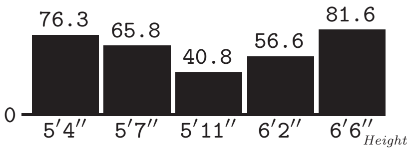

Fig. 7. % of 'False' responses to 'X is tall and not tall'

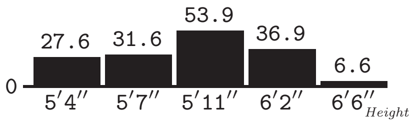

Fig. 8. % of 'True' responses to 'X is neither tall nor not tall'

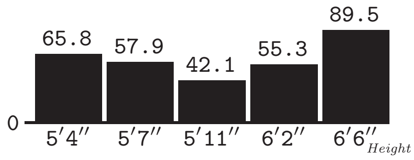

Fig. 9. % of 'False' responses to 'X is neither tall nor not tall'

relevant question: T_b, for example, is the number of truth-judgers for '#2 is tall and not tall'; F_n is the number of falsity-judgers for '#2 is neither tall nor not tall'.[13] What we want to highlight is that *neither*, whose truth can justify a truth-value gap, coincides in many cases (more than half!) of borderline-height with *both*, which, when true, suggests a truth-value glut.

Another interesting correlation is the one found between the questions 'x is tall' and 'x is not tall' on the one hand, and 'x is tall and not tall' on the other. Figure 10 shows that 32.4% of those who thought it was true that #2 was 'tall and not tall' also thought it was *false* that he was tall and *false* that he was not tall. Figure 11 illustrates the correlation in the other direction; it shows the percentage of *true* responses to 'x is tall and not tall' when the statements 'x is tall' and 'x is not tall' are judged false. The ratio is 68.8% at 5′11″, and 100% at 6′2″.[14]

There are other interesting findings that center around our borderline suspect #2. For example, we find the number of subjects who think '#2 is tall' and '#2 is not tall' are both true to be much higher than those who think they are both false (Table 5).

13 A Bowker's test for symmetry gives an X^2 value of 8.04. With $df = 3$, the tail probability $p < 0.05$.

14 We think the anomalous value of 100% for our 6′2″ subject is due to the fact that only four subjects thought that both 'tall' and 'not tall' were false for this subject. And all of these four thought #5 was 'tall and not tall'.

Table 4. Cross-tabulation of 'neither' and 'both' (Height = 5′11″): Response types for 'tall and not tall' are subscripted with b (for 'both'), and those for 'neither tall nor not tall' are subscripted with n. For example, the number of truth-judgers for both questions is in the cell where the T_b row intersects with the T_n column.

	T_n	F_n	C_n	
T_b	22	12	0	34
F_b	13	18	0	31
C_b	6	2	3	11
	41	32	3	76

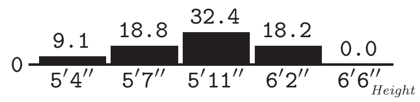

Fig. 10. % of Falsity of 'tall' and 'not tall' when *both* is true

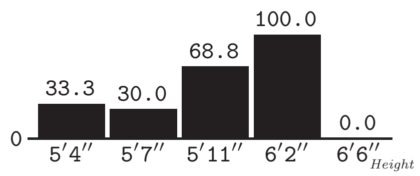

Fig. 11. % of Truth of *both* when 'tall' and 'not tall' are false

5.2 Analysis and Implications

Our goal in this section is to suggest a possible explanation for the pattern that we have just demonstrated: the pattern where 'is tall' and 'is not tall' are both considered false (when they are about a borderline individual), but where 'is tall and not tall' and 'is neither tall nor not tall' are considered true of that same individual.

Our idea, as we promised, relies crucially on the Gricean maxims of conversation. However, the solution also relies on an assumption that may seem somewhat controversial: that a given vague predicate has two possible interpretations, a *super*-interpretation and a *sub*-interpretation, in the same way that a vague expression containing negation can be interpreted strongly (i.e. super-interpreted), or weakly (i.e. sub-interpreted). Assuming this, together with the Gricean maxims, provides a way of accounting for the seemingly inconsistent patterns outlined above.

Our intuitive semantic theory is rather standard and classical, so far as our inclusion of vague and ambiguous predicates allows. Given a domain $\mathcal{D}$ of individuals, the extension of a non-vague predicate[15] is interpreted normally, as some subset of $\mathcal{D}$ (the ones that manifest the property). The negation of a predicate is simply its complement, relative to $\mathcal{D}$ (note that we are assuming an unambiguous definition of negation here). A predicate that is vague but not ambiguous is represented as a set of ordered pairs, the first member of which is an (admissible, classical) precisification of the predicate and the second is the subset of $\mathcal{D}$ that satisfy the predicate in that precisification. The extension of that predicate is always relative to one or a group of the precisifications, and then becomes the subset of $\mathcal{D}$ that obeys that restriction on the precisifications. A (two-way) ambiguous predicate is interpreted as having two members, each one of which is an interpretation of the former types. One kind of meaning that a vague predicate can manifest is what we have intuitively called 'the super-interpretation': its extension is the subset of $\mathcal{D}$ of things that occur as values in *every* precisification. Another is what we called 'the sub-interpretation': the subset of $\mathcal{D}$ that appear as values of some precisification.

Table 5. Percent of Ss who gave same answer to both '#2 is tall' and '#2 is not tall'

#2 is tall	#2 is not tall	percent
T	*T*	3.9%
F	*F*	21.0%
C	*C*	5.3%

Our view here is that a vague predicate such as 'tall' can be ambiguous between the super- and sub-interpretations. A hearer is to find the more suitable interpretation of the predicate from these two possible meanings, or just to say that there is no way to choose and the sentence is simply ambiguous. The Gricean Maxim of Quantity can then be a condition on which one of these sets should be selected from the interpretation of the predicate. The condition is that the selected set may not be a superset of any other member of the set.[16] For 'tall', the result will be the set of the super-tall people, since it is the only set, from the two available options, for which the condition holds. For 'not tall', the two possibilities are the complements of the super-tall set and the sub-tall set, since 'not' unambiguously denotes the set-complement operation in this conception (complement with respect to $\mathcal{D}$).

So, the set of available interpretations will contain both the complement of the super-tall individuals, and the complement of the sub-tall individuals. The former set, the complement of the set of super-tall individuals, is the set of the individuals that are not super-tall, i.e., the borderline cases and the definitely not-tall cases. The latter set, the complement of the set of sub-tall individuals,

[15] We would extend this to n-place relations, but for the present paper we stick to monadic predicates.

[16] We see this as an application of the Strongest Meaning Hypothesis of Dalrymple et al. [3].

will contain individuals that are not sub-tall, that is, every individual *except* those that belong to the extension of tall in some sharpening. In other words, the set will contain the individuals for whom there are *no* sharpenings in which they belong to the extension of tall, and since each sharpening is classical, they are precisely the individuals that are super-not-tall. (Another name for this set might be the super-short individuals). So, the two possible interpretations for 'not tall' are the set of borderline-cases together with the super-not-tall cases (from the complement of the super-tall set of individuals), and the set of super-not-tall individuals (the complement of the sub-tall individuals, the super-short ones). And according to the condition of quantity, the latter set is selected since there are no subsets of itself that belong to the collection of interpretations. Formulated this way, the Maxim of Quantity will favor the super-interpretation both for 'tall' and for 'not tall', making it seem that negation is choice negation when it is really the effect of these pragmatic operations.

We now show how 'tall and not tall' might be made to mean 'borderline' on this approach. The set of interpretations will contain four elements, each of which resulting from the intersection of two sets (through the denotation of 'and'). The four elements are shown in (3). Note that, of the four options, only (3c) can be nonempty.

(3) a. $\{x : x \text{ is supertall}\} \cap \{x : x \text{ is supertall}\}^- = \emptyset$
b. $\{x : x \text{ is supertall}\} \cap \{x : x \text{ is subtall}\}^- = \emptyset$
c. $\{x : x \text{ is subtall}\} \cap \{x : x \text{ is supertall}\}^- \neq \emptyset$
d. $\{x : x \text{ is subtall}\} \cap \{x : x \text{ is subtall}\}^- = \emptyset$

(3a) and (3d) are empty because in both cases the intersection is applying to a set and its complement. (3b) is empty because the individuals it will contain are those that are tall in every precisification and at the same time not tall in any. Now, if in the formulation of the maxim of Quality one can block readings that are trivially false, then the maxim will allow only (3c) to emerge from the four options in (3). 'Tall and not tall' can therefore only denote the set of individuals who are sub-tall and sub-not-tall, i.e. the borderline individuals.[17]

Finally, in the case of 'not not tall', there are also two available interpretations: for 'tall' as super-tall we get the complement operation canceling itself, by applying twice, and yielding the set of super-tall individuals, and likewise, for 'tall' as sub-tall, we get the set of sub-tall individuals. Of the two options, it is the set of super-tall individuals that will qualify, and so we predict, incorrectly, that 'not not tall' means super-tall. But here we may add that the Gricean maxim of Manner, which penalizes prolixity, will block the super-interpretation, for if the super-interpretation was intended, the speaker would have had no reason to use 'not not' in his/her locution, but rather would say simply 'is tall'. It

[17] This is where intuitionistic negation fails to make the correct prediction: the negated conjunct is negated only once, so it can only be interpreted strongly. But in order for the conjunction to denote a non-empty set, the negation has to be interpreted weakly.

is difficult to precisely formulate a mechanism that blocks candidate interpretations on the basis of brevity, but as the maxim has generally proven useful in the theory of pragmatics, we feel it innocuous to invoke it for our purposes, hoping that however it can be made formal, it can be utilized to disqualify the super-interpretation from entering the set of denotations for 'not not tall', and instead interpreting the predicate as sub-tall.

We wish to emphasize that our theory does not suddenly use both super- and sub-valuationist interpretations in an ad hoc manner merely for the special case of apparently contradictory statements. Indeed, the assumption that they are both in play is not specific to borderline cases at all. Rather, we suggest that the interpretation of a vague predicate, regardless of whether or not the property is being predicated of a borderline individual, can be modeled using sets of precisifications, which is the architecture that both super- and sub-valuationists share. Where the two approaches diverge is in the use of the quantifier; in supervaluations, p is true when p holds in *all* precisifications, while in subvaluations, p is true when it holds in at least one precisification. What we suggest is that the use of the quantifier is pragmatically governed. Informativity (that is, Gricean quantity) demands the stronger of the two quantifiers, i.e., the supervaluationary interpretation, but in the case of contradictory statements like '#2 is tall and not tall', using the universal quantifier produces a trivially false statement; so the quantifier must be weakened in order to make the statement non-trivial, and we propose that it is weakened to an existential quantifier, thereby producing the subvaluationary interpretation.

There is a possible view – though not well-motivated, we hope to show – according to which our patterns are interpreted as support to the fuzzy approach to vague expressions. Recall that in fuzzy logic there is an infinite number of truth values, ranging from 0 (false) to 1 (true), and that the truth-value of $\neg p$ for any proposition p is $1 - V(p)$. Thus, for example, if $V(p) = 0.6$, the value of its negation $\neg p$ is $1 - 0.6 = 0.4$. Recall also that the truth value of a conjunction $p \wedge q$ is defined as the minimum of the truth values of the conjuncts p and q. If the truth-value of p were 0.6, for example, and the value of q were 0.3, then the value of $p \wedge q$ will be $min(p, q) = 0.3$. This makes it possible for contradictory expressions like $p \wedge \neg p$ to be more true than 0; for if the truth-value of p were 0.6, the value of $\neg p$ will be 0.4, and the value of the conjunction $p \wedge \neg p$ will be $min(0.6, 0.4) = 0.4$.

A fuzzy logician may point to Figs. 6 and 7 and claim that the findings they illustrate are in fact faithful to the predictions of fuzzy logic, specifically, the prediction that a contradictory proposition containing a vague predicate is false at the periphery, and gradually climbs to half-truth in borderline cases. The same could be said to hold with respect Figs. 8 and 9, if the disjunction of p and q is computed as $max(p, q)$. A defender of this view may add that the patterns in Figs. 2-5 lend further support, since the truth of relevant propositions seem to gradually climb from near-falsity on one end of the tallness spectrum, to near-truth on the other end.

The problem with this view is that it assumes a statistical notion of truth, that is, a definition of truth whereby a proposition is said to be true to a degree determined by consensus. We think that proponents of this view argue in favor of the fuzzy approach without taking notice of how believers of contradictions – the truth-judgers of 'tall and not tall' – judge the truth of other related statements like 'x is tall' and 'x is not tall'. In other words, while the *percentages* of truth/falsity-judgements made by many different people can indeed be thought to resemble a fuzzy pattern, a closer look at how the same judgers, taken individually, responded to other queries reveals a recurrent pattern that the fuzzy approach cannot predict, namely, the pattern in which a borderline proposition, and its negation, are judged false, but in which their conjunction is simultaneously judged true.[18]

6 Conclusion

We have argued that the findings of BOVW were incorrectly interpreted as support for the VAGUENESS-AS-IGNORANCE hypothesis. In the course of our argument we suggested that BOVW's theoretical criticisms against the gap-theoretic account of higher-order vagueness are inconsistent with their defense of their own proposal. We also showed that BOVW question-beggingly presuppose a bivalent proof system in their claim that gap-theories lead to contradictory statements, and also that their experimental evidence for the logical equivalence of 'x is not tall' and '"x is tall" is false' was not convincing. Finally, we presented new experimental findings that contradict BOVW's explanation of gaps: the emergence of gaps, they claim, is due to a general preference for errors of omission. If this claim were valid, we would expect a much larger percentage of 'Can't tell' responses in borderline cases. This, however, was not the case.

We ended our discussion by shedding experimental light on a different view of vagueness, a view in which a predicate and its negation are each said to be false of a borderline individual, but in which their conjunction is said to be true. We acknowledge, however, that further experimental work (as well as further theoretical work) is called for to test the details of super- and sub-valuations and their interactions with the Gricean maxims. Our 'pragmatic story' is but a first step which needs further investigation.

References

1. Alxatib, S., Pelletier, F.J.: The psychology of vagueness: Borderline cases and contradictions. Mind and Language (forthcoming)
2. Bonini, N., Osherson, D., Viale, R., Williamson, T.: On the psychology of vague predicates. Mind & Language 14, 377–393 (1999)
3. Dalrymple, M., Kanazawa, M., Kim, Y., McHombo, S., Peters, S.: Reciprocal expressions and the concept of reciprocity. Linguistics and Philosophy 21, 159–210 (1998)

[18] For further criticism of the fuzzy account of vague predicates from an experimental point of view, see Ripley [14]. Note particularly his finding that subjects tend to *fully* agree with (allegedly) contradictory statements – choosing 7, 'Agree', on a scale of 1-7, rather than choosing a more moderate response, as the fuzzy logician would predict.

4. Fine, K.: Vagueness, truth and logic. Synthèse 30, 265–300 (1975)
5. van Fraassen, B.: Singular terms, truth-value gaps, and free logic. Journal of Philosophy 63, 481–495 (1966)
6. Horn, L.: A Natural History of Negation. University of Chicago Press, Chicago (1989)
7. Hyde, D.: From heaps and gaps to heaps of gluts. Mind 106, 641–660 (1997)
8. Jaśkowski, S.: A propositional calculus for inconsistent deductive systems. Studia Societatis Scientiarum Torunensis 1, 55–77 (1948); Article published in Polish with the title "Rachunek zdań dla systemów dedukcyjnych sprzecznych." English translation in Studia Logica 24, 143–157 (1969)
9. Kamp, H.: Two theories about adjectives. In: Keenan, E. (ed.) Formal Semantics of Natural Language. Cambridge University Press, Cambridge (1975)
10. Keefe, R.: Theories of Vagueness. Cambridge University Press, Cambridge (2000)
11. Levinson, S.: Presumptive Meanings: The Theory of Generalized Conversational Implicature. MIT Press, Cambridge (2000)
12. Machina, K.F.: Truth, Belief, and Vagueness. Journal of Philosophical Logic 5(47-78)
13. Pelletier, F.J., Stainton, R.: On 'the denial of bivalence is absurd'. Australasian Journal of Philosophy 81, 369–382 (2003)
14. Ripley, D.: Contradictions at the borders. Paper Presented at the 2008 Conference on Philosophy and Psychology of Vagueness at the Institut Jean-Nicod (2008)
15. Ritov, I., Baron, J.: Reluctance to vaccinate: Omission bias and ambiguity. Journal of Behavioral Decision Making 3, 263–277 (1990)
16. Sorensen, R.: Blindspots. Clarendon Press, Oxford (1988)
17. Sorensen, R.: Vagueness and Contradiction. Oxford UP, New York (2001)
18. Spranca, M., Minsk, E., Baron, J.: Omission and commission in judgment and choice. Journal of Experimental Social Psychology 27, 76–105 (1991)
19. Williamson, T.: Vagueness. Routledge, Oxford (1994)