The Psychology of Vagueness: Borderline Cases and Contradictions*

Sam Alxatib
MIT

Francis Jeffry Pelletier
University of Alberta

Abstract

In an interesting experimental study, Bonini et al. (1999) present partial support for truth-gap theories of vagueness. We say this despite their claim to find theoretical and empirical reasons to *dismiss* gap theories and despite the fact that they favor an alternative, epistemic account, which they call 'vagueness as ignorance'. We present yet more experimental evidence that supports gap theories, and argue for a semantic/pragmatic alternative that unifies the gappy supervaluationary approach together with its glutty relative, the subvaluationary approach.

Word Count (approximate, using DETEX | WC): 9355.

1 Introduction

The history of philosophy has seen 'the problem of vagueness' raised as an ontological question (concerning whether reality can be vague), a logical question (about how to reason consistently using vague terms), an epistemological question (covering issues of how we can ever know anything, given the existence of vagueness), a linguistic question (about how to describe the meaning of vague terms) and a conceptual question (about how it is possible to control one's views about reality that use vague concepts). Historically, the problem of vagueness first arose in its logical guise, when in the 4th century BCE, Eubulides of Miletus formulated what is known today as the Sorites Paradox, the paradox of the heap. The paradox results from induction on premises like the following:

(1) 100,000 grains of wheat make a heap.

(2) if n grains of wheat make a heap, then $n-1$ grains of wheat make a heap.

The combination of assumptions (1) and (2) leads to the conclusion that one (or even zero) grain(s) of wheat make a heap.

Eubulides' contribution to the problem of vagueness focused on the logical issue, concerning how we should represent and reason with vague concepts. Each premise seems intuitively to be

*Address for correspondence: F. J. Pelletier, Dept. Philosophy, Univ. Alberta, Edmonton, AB, Canada T6G 2E7
email: Jeff.Pelletier@ualberta.ca

true and unobjectionable, and yet the conclusion is clearly false. In the course of trying to discover what is wrong with Eubulides' argument, different theorists have focused on different aspects of the general topic of vagueness, some focusing on our knowledge of vague concepts/meanings, others on such ontological issues as whether there can be vague-objects-in-reality, and others on the role of representing vagueness in an artificial language designed to show how reasoning with vague representations is possible. The full story of vagueness will doubtless require a coherent theory encompassing all these aspects.

Most current discussions of vagueness pay particular attention to the notion of a borderline object: objects to which the purportedly vague term neither does apply nor does not apply—at least, not without some further considerations. A person with some intermediate number of hairs on his head[1] is not bald, nor is he not not-bald—at least, not without further qualifications. (Of course, these further qualifications can take many forms, including the view that, although this is what vagueness *would* be, there is in fact no such thing as vagueness in reality—only our inability to *know and say* whether or not these so-called borderline cases are cases of bald or of not-bald.)

Much of the current discussion of vagueness presumes that an adequate solution to Eubulides' puzzle will yield more general answers to the broader topics within 'the problem of vagueness'. So, in Bonini et al. (1999), henceforth BOVW, an examination is conducted of the way that vague concepts are viewed by people who might then use them in Sorites-like arguments. And they purport to discover that the solution concerns the way such concepts are unconsciously interpreted. The solution they favor is called an 'epistemic account', but to appreciate their position it is necessary to survey some competitor accounts of ways to solve the Sorites Paradox.

1.1 Solutions

Following BOVW's lead, we will not enter into the substantial topic of defining the underlying cognitive structures of vagueness in terms of concepts and prototypes (e.g., Hampton, 2007; Hampton et al., 2006), nor do we (or BOVW) discuss how vagueness might be represented in formal semantics (Barker, 2002; Frazier et al., 2008; Kennedy, 2007). Instead, we focus on the logical side of the topic. Of the numerous proposed solutions to the logical problem of the sorites, only three pertain to our current discussion: the method of supervaluations (a 'gap' theory), that of subvaluations (a 'glut' theory), and the epistemic solution (which purports to be a 'classical logic' theory). Other possible solutions, such as three-valued logics (Tye, 1994), fuzzy logic (Zadeh, 1975; Smith, 2009), context-dependent approaches (Graff, 2000; Stanley, 2003; Shapiro, 2006), modal approaches (Pelletier, 1984; Bennett, 1998), bi-modal logics (Halpern, 2008), default logics (Cohen et al., 2008), and conversational policy (Scharp, 2005; Walton, 2006) are not dealt with directly by BOVW and will not be directly discussed here, except in passing. We describe the three approaches that we

[1] Distributed in a certain way, of course. But we won't bring this aspect up in what follows.

and BOVW are concerned with – rather informally, but enough for our purposes – in the following four subsections.

1.2 Supervaluations

The idea behind the supervaluationary method comes from Mehlberg (1958), but was first formalized in van Fraassen (1966), and later elaborated in the context of vagueness by Fine (1975); Kamp (1975) and Varzi (2007). The idea is to associate a given vague predicate with mutiple 'sharpenings' (also called precisifications), each of which contains some precise cut-off point. Within each sharpening, the individuals whose membership to P matches or exceeds the cut-off can be said to belong to the extension of P *in that precisification*. Because each precisification is classically constructed, the individuals that do *not* belong to P are in P's negative (or anti-) extension *in that precisification*. The predicate P is said to hold, without qualification, of an individual a if and only if Pa holds in every one of these sharpenings.

It is easiest to see this with an example that invokes an underlying ordered scale, such as a scale of height. So, for example, suppose the predicate *tall* is assigned a (highly restricted) collection of just three precisifications, one of which says that the cut-off is 180 cm, another marking it at 182 cm, and another at 177 cm. If person a is 185 cm, then a will belong to the extension of *tall* in every one of our three ways of making *tall* precise. In this case, the statement 'a is tall' is said to be *super*true (assuming there to be no other precisifications). In supervaluations, a statement is considered true (without a restriction to a sharpening) if and only if it is supertrue. This also makes the statement 'a is not tall' be superfalse and hence false. Now suppose person b is 170 cm. In each precisification, b belongs to the anti-extension of *tall*, which makes the statement 'b is tall' superfalse, and therefore false. This also makes the statement 'b is not tall' supertrue, thus true. Now, if person c is 179 cm, c will belong to the extension of *tall* in some precisifications, and its anti-extension in others. Thus the statement 'c is tall' is true in some of these precisifications, but not in others. So it cannot be considered true, but it cannot be considered false either, since it is neither supertrue nor superfalse. This case demonstrates how truth-value gaps arise: if the precisifications are such as we are imagining, 'c is tall' is not assigned either truth-value because it is neither supertrue nor superfalse, hence neither true nor false. There is a gap in truth value, between true and false.

Let us now reconsider assumption (2) above (the inductive part of the sorites). This (quantified) conditional says that, for any value n, if n grains of wheat make a heap, then $n-1$ grains of wheat make a heap. This is not true in a supervaluationary setting: because every precisification is *classical*, so that in every precisification there will be some distinct value m such that m grains of wheat make a heap, and $m-1$ grains of wheat do not make a heap. This makes premise (2) false in every precisification, i.e. superfalse, thus defusing the sorites argument. Note that m differs among the precisifications. So while it is supertrue for a predicate P that there is an m for which $P(m) \wedge \neg P(m-1)$, the value m cannot be specified because it varies from sharpening to sharpening. It is this feature that is said to give the paradoxical feel to the Sorites Paradox, according to the supervaluationist.

1.3 Subvaluations

The method of subvaluations shares its basic architecture with supervaluations: a predicate P is multiply precisified, and each precisification is classically constructed. But the crucial difference is that, here, a statement like Pa is said to be true if and only if it is *sub*true, i.e., just in case Pa holds in *at least one* precisification. (Falsity, likewise, requires subfalsity.) The framework finds it origins in the discursive logic of Jaśkowski (1948), and is later formulated for the case of vagueness by Hyde (1997). To illustrate the consequences of construing truth as subtruth, let's look back at our simplified example above: *tall* is precisified in three ways: once sharpened at 182 cm, again at 180 cm, and again at 177 cm. For person a, whose height is 185 cm, the statement 'a is tall' is true, since there is a sharpening – in fact there are three – in which a exceeds the minimum for tallness. Furthermore, because there is not a single sharpening in which 'a is tall' is false, the statement 'a is tall' is not false. So it is True-And-Not-False. On the other hand, person b, who is 170 cm high, does not belong to the extension of *tall* in *any* of the given precisifications. Therefore, the statement 'b is tall' is not true (because it is not subtrue); but it is false, because it is subfalse. So, it is False-And-Not-True. Once again, the interesting case is that of person c, our borderline example. There is a sharpening in which c is tall, and there is also a sharpening in which c is not tall. The statement 'c is tall' is therefore *Both*-True-*and*-False; and for the same reason 'c is not tall' is Both-True-And-False. Assigning multiple truth values to a single proposition creates what is called a truth-value *glut*.

As in the supervaluation case, for a predicate P, it is true that there is an m for which $P(m) \wedge \neg P(m-1)$, since this holds in every precisification; and hence this True-And-Not-False. However, since the value m varies from sharpening to sharpening, for each one of those values it will be both true and false that they mark the cutoff point. A holder of the subvaluation theory will attribute the paradoxical nature of the Sorites Paradox to this feature.

1.4 Epistemic Theories

Epistemological theories of vagueness are so-called because they attribute vagueness to a lack of knowledge. Such theories deny premises like (2) above, insisting that *in reality* there in fact *is* some particular m such that having $m+1$ grains of sand (piled atop one another) makes a heap while having just m grains similarly piled does not. Sorensen (1988; 2002) and Williamson (1994) are the two main modern advocates for this explanation of the source of vagueness. The view is rather like taking one of the supervaluationist's precisifications and proclaiming it to be the 'correct' one — the one that correctly characterizes reality. The vagueness arises because we don't *know* this m. In fact, according to most vagueness-as-epistemological theorists, we *can't* know the value. After all, they say, by hypothesis of the sorites argument, there is no difference in the evidence available to a person looking at a heap with m grains vs. looking at one with $m+1$ grains. One can't *know* that the pile with m grains is a heap, since it is in fact *not* a heap, according to their theory. But looking at one with $m+1$ grains gives the very same evidence available to the viewer, so the viewer can't *know* that it *is* a heap, even though it is. It is this fact, according to the epistemic theorists, that gives the Sorites Paradox its paradoxical nature.

1.5 Some Logical Features of the Three Theories

It is a natural reaction to the phenomenon of vagueness to suggest that the logical representation system should have three values: true, false, and vague. However, such three-valued logics require the theorist to make some seemingly arbitrary choices and contain some quite unusual properties. For instance, with a third value, it becomes a matter of choice whether 'if p then p' should be a logical truth or not. If the theorist chooses not, then the representation system foregoes the well-established deduction theorem: it would become possible for there to be valid arguments of the form 'from premises A, B, and C we can correctly infer D' but where it would be an invalid argument to say 'from premises A and B we can infer "if C then D" '.[2] On the other hand, if the theorist chooses to say that it *is* a theorem, then 'if p then not-p' would be true when p takes the third value.[3] And that also violates common usage.

Since classical first order logic is the best-understood representational system, theorists naturally gravitate toward it when they are trying to capture novel features of language. First order logic has a complete and sound deductive system associated with it, there is a clear theory of truth-in-a-model for it, and, generally, we know all the logical features of such a system. Adopting this system would secure a firm footing for any novel feature that can use it. Theorists hoping to capture vagueness in the theory of language tend therefore to demonstrate that their approach is, or at least is compatible with, classical first order logic.

Most straightforwardly, the epistemic theory adopts this mode of representation directly. Vagueness, they say, is not a feature to be represented in the theory, since it does not exist except in the minds of the users of the language. And they deride theorists who move away from classical logic. But keeping the desirability of first order logic in mind, some 'gap theorists' have been able to show that it is possible to keep much of classical logic while admitting gaps nonetheless. This is the

[2] For example, from the premise p we can infer p (because every interpretation that makes the premise true also makes the conclusion true), but 'if p then p' would not be logically true, since it would not be true when p takes the third value.

[3] Because the proposed theoremhood of 'if p then p' requires that even when p assumes the third value this sentence is true. But if p has the third value, then so does not-p, and hence 'If p then not-p' would be true too.

ploy taken by supervaluationists: since all of the logical truths of first order logic are true in every precisification (naturally, since after all, each sharpening is classical), it follows that every logical truth of classical logic is also a truth of supervaluation theory. And therefore it too is a 'logically conservative' representational language. A difference is that supervaluation theorists allow some sentences to be neither true nor false; but none of these are truths of classical logic.

Subvaluation theories seem to be at odds with classical logic, since they allow statements to be *both* true and false. But there are some developments of this theory that approach classical logic, such as the logic LP of Priest (2006) and the theory of Jaśkowski (1948) as developed in Hyde (1997). The theory of inference developed in these theories seems just as robust as those developed for supervaluatons, and so it is hard to see that is a strong argument to claim that these theories are deductively inferior to classical logic.

2 BOVW's Experiment

2.1 Method

The experimental evidence of BOVW was gathered by means of questionnaires given (in Italian) to a total of 652 students at Italian universities. 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. For the predicate *tall*, for example, they provided the instructions seen below. (The queries regarding 'truth' were given to a different group of subjects from the ones regarding 'falsity'. There were 320 'truth-judgers' and 332 'falsity-judgers' in total).

> 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.

> 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.

The same design was used to elicit responses for *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 a set of variant trials, which will be important in our discussion, the words 'true' and 'false' were removed from the query, and the instructions were modified to the following:

> 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.
>
> 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.

2.2 Results

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, it was found 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 former are called 'truth judgments' and the latter 'falsity judgments' by BOVW.) In four of their six trials, the results are in Table 1.[4]

	Trial 1	Trial 2	Trial 4	Trial 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

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

These findings contradict the predictions of glut-theories of vagueness. As a case in point, consider Hyde's subvaluationary framework. Here, a statement such as 'a is tall' is considered true just in case there exists at least one precisification of *tall* in which a belongs to the extension of the sharp version of *tall*. Given a collection of admissible precisifications, an individual a of borderline height will belong to the extension of *tall* in some precisifications, and to its anti-extension in some other precisifications, and because of the 'weak' requirement on truth in the framework, 'a is tall' will turn out true *and* false simultaneously. In other words, 'tall' is considered true of every individual ranging in height from the very tall down to the lower end of the borderline range, and

[4] The predicate 'tall' was not used in their Trial 3. Trial 6 was somewhat different from the other five trials because subjects were explicitly alerted to the existence of 'middle ranges' of values. (For example, in place of the most recently quoted variation, subjects were told 'When referring to the height of a person, we can distinguish between "tall", "medium-height", and "short". Of course, "tall" applies to people of great height and not to those who are short or of medium height.') This sort of priming might alter the subjects' views about the relation between 'is not tall' and 'is false that he is tall'. And in fact, the results for various of the tested items seem quite different in this trial than they do in other trials. So we have decided not to include the data from Trial 6.

'not tall' is considered true of everyone from the very not-tall up to higher end of the borderline range. (One could also say that 'tall' is *false* of every individual from the very bottom of the spectrum to the upper edge of the borderline). This is illustrated in Figure 1. The answer to the 'false' questions is therefore predicted to be higher than the answer to the 'true' questions, which is the opposite of what was found in BOVW.

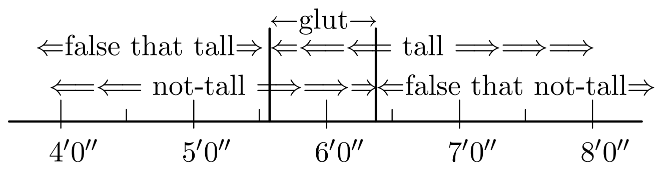

Figure 1: Subvaluation predictions about height

Gap theories such as supervaluations, on the other hand, find support in these results. Since truth is supertruth, that is, truth in *every* admissible precisification, 'x is tall' will hold only of those xs that are tall in every way of making the predicate *tall* precise. Similarly, 'x is not tall' will be true just in case x is not tall in every precisification. An individual a of borderline height is therefore neither tall nor not tall, since neither statement holds true of a in every precisification. The prediction, then, is that the minimum height for tallness lie at the higher end of the borderline range, and that the maximum height for not-tallness lie at its lower end. The former is thus expected to be higher than the latter, as was found experimentally by BOVW. This is illustrated in Figure 2.

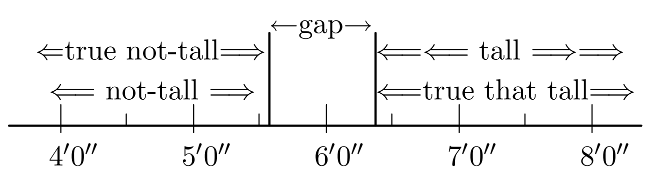

Figure 2: Supervaluation predictions about height

2.3 Evaluation

Surprisingly, however, BOVW reject the gap account and instead promote the following epistemic hypothesis:

(5) 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.

Gaps appear, according to them, because speakers are more willing to commit errors of omission than commit errors of commission, that is, they 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.[5] 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 they 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.

Re (1): 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 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.

The treatment of higher-order vagueness varies across theories, but in supervaluations a possible maneuver is to allow borderlineness to apply not only to the predicate in question, but also to the admissibility of the way the predicate is made precise (this is informally sketched in Keefe, 2000). Suppose that the predicate *tall* is sharpened in multiple ways, where each sharpening consists of a distinct and precise cut-off point. If the truth of a given statement depends on its truth in every admissible precisification, its valuation can only produce one of three crisp possibilities: either it is true, false, or neither. But, clearly, the boundaries are not so easily delineable. To create more gradations, it is added that some precisifications are admissible, some are not, and some are neither. Imposing a cut-off for tallness at 200 cm is certainly not admissible, and the corresponding precisification is therefore not considered when assessing the tallness of a given individual.[6] Similarly, 160 cm is too low and is also not admissible. But somewhere in between there can be several admissible sharpenings. If admissibility is made vague, then somewhere between

[5] Based on studies by Ritov and Baron (1990); Spranca et al. (1991).

[6] Here we are considering a comparison class where the average height is near 180 cm. That is, we are excluding from this illustration classes like basketball players and little people, to whom our numeric examples may be more controversial.

200 cm and 180 cm, say, there will be some precisification, call it s, which is neither admissible nor inadmissible. Now suppose that a's height is just below the cut-off point in s. a's tallness depends on whether or not s is admissible: if it is, then a will be tall in some, but not all, admissible precisifications, making him borderline; if s is inadmissible, then a will be tall since s will not contribute to the computation at all; and if s is neither admissible nor inadmissible, then a will be neither tall nor borderline, that is, a will be considered 'borderline-borderline-tall'. Note that the notion of admissibility is a *metalinguistic* notion. It follows, then, that a higher metalanguage (a meta-metalanguage) is needed in order to describe the semantics of admissibility in the lower metalanguage. The meta-meta-linguistic conditions on admissibility may also be susceptible to borderlineness, a feature which would then require yet another, higher metalanguage, and so on. Theoretically, the process can be repeated indefinitely, since the 'standards' imposed in every metalanguage could be vague. So, in effect, the finer gradations are realized by alluding to higher levels of borderlineness.

The problem with this approach and the reason they reject supervaluation theory, according to BOVW (p. 388), is that 'the mental representation of all these vague boundaries seems psychologically implausible'. 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 they use introspection as a test of acceptability of a semantic theory. We note, as they do, that there is also 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.

Re (2): The second argument of BOVW against gap-theories starts with the claim that their studies show that gap-theorists must be committed to the position that the statements 'n is tall' and ' "n is tall" is true' have the same truth conditions (ditto for 'n is not tall' and ' "n is tall" is false'). In the gap-theorist's case, if she were to deny this, then she must expect wider gaps to emerge for the metalinguistic statements because, in a framework like supervaluations, truth and falsity are associated with supertruth and superfalsity. A statement is supertrue if and only if it is true in every precisification. So, ' "n is tall" is true' if and only if 'n is tall' holds in every way of making *tall* precise. If the truth conditions for 'n is tall' were different from the truth-conditions for ' "n is tall" is true', then the expectation is that the former's truth-conditions should be more lenient, so to speak, than the conditions for the latter. This means that the minimum n that makes it *true* to say that 'n is tall' should come out greater than the minimum n which makes n tall, and the maximum n that makes it *false* to say that 'n is tall' is expected to be lower (i.e. more restricted) than the maximum n which makes n not tall. In other words, the metalinguistic statements should produce bigger gaps than the non-metalinguistic ones.

The empirical evidence BOVW present, however, is alleged to disprove this prediction, as they claim no difference was found between the metalinguistic gaps and the non-metalinguistic ones. For *tall*, the average answer of truth-judgers in the first metalinguistic trial (their Study 1) was 178.30 cm, and in the second (Study 2), the average was 179.55 cm. In comparison, the average answer of truth-judgers in the first *non*-metalinguistic trial was 181.49 cm (Study 4), and 178.28 cm in the second (Study 5). Conversely, the falsity-judgers gave higher values – less strict, that is – for the metalinguistic questionnaires. In Study 1 (metalinguistic), the average was 167.22 cm, and in the second trial (Study 2) the average was 164.13 cm. In Studies 4 and 5 (non-metalinguistic), the average answers for falsity-judgers were 160.48 cm and 163.40 cm, respectively. Below we will challenge this interpretation, but for now we continue with their line of argumentation.

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 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 to show that this cannot be correct:

(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 must therefore be revised. But the only assumptions concerned the existence of borderline gaps that are posited by supervaluationists, together with the use of logic in moving from one statement to another.

Gap theorists generally, and supervaluationists in particular, will find this general form of argumentation unsatisfactory. As can be seen from the argument, there is the move to link up metalinguistic statements of the form ' "n is tall" is true' with statements like 'n is tall'. The former, metalinguistic statement predicates 'is true' of an object language statement. In the argument, the unquoted occurrences of statements like 'n is tall' are the metalinguistic translations of the quoted versions of the same statements. Thus the argument, carried out in the metalanguage, is trying to link up semantic predications made of object language sentences with some other metalinguistic statement. The linkage is most explicit in argument steps (3)–(7), all of which are dependent on the claims made in steps (3) and (4) together with the logical moves involving negation used to go from these to steps (5)–(7).

We think that a supervaluationist could legitimately complain about the inferences made using negation. Suppose we agree with (4) and (5). It is known that, in supervaluations, a statement S is false if and only if it is false in every admissible precisification. The premise in (4) says that the same applies to 'not S'. That is, 'not S' holds just in case 'not S' holds in every admissible precisification. Supervaluationists also say that S is true if and only if it is supertrue. So, if S is not supertrue, S is not true. In other words, the conditions under which 'S is not true' holds are those that make S *false or truth-value-less.* According to the supervaluationist, then, (4) says that 'not S' holds just in case S is false, and (5) says that 'not S' holds just in case S is false or truth-value-less. Because the conditions for falsity and 'not-truth' are different, premises (4) and (5) can only be simultaneously maintained if the negation in 'not S' is treated ambiguously. In (4), 'not' denotes what is popularly known as choice negation, and in (5) it denotes exclusion negation. From this distinction it follows that (6) does not lead to (7), and the contradiction in (10) cannot be deduced.

Put another way, a gap theorist will find that the argument presupposes a bivalent logic in its treatment of negation. If we do not make this supposition, but we still agree with the experimental result reported in (3), we would then say that every set of circumstances in which x is tall is also a set of circumstances in which it is true that x is tall, and vice versa, i.e., that

(a) $p \vDash \text{True}(p)$

(b) $\text{True}(p) \vDash p$

are both correct, and that they justify

(c) not-$p \vDash$ not-True(p) i.e., not-$p \vDash$ (either False(p) or Gap(p))

(d) not-True(p) $\vDash$ not-p i.e., (either False(p) or Gap(p)) $\vDash$not-p

The Gap-theorist agrees with the Epistemicist that

(e) False(p) $\vDash$not-True(p)

and so

(f) False(p) $\vDash$not-p

But the Gap theorist does not agree that this last implication can be reversed, for the reason announced in the explanation of (c)—that from not-p one arrives at the *disjunction*, either False(p) or Gap(p). So the Gap theorist believes that, in general,

(g) not-$p \nvDash$ False(p),

and therefore thinks that step (4) of BOVW's argument,

(4) n is not tall $\equiv$ 'n is tall' is false

is wrong, because that it illegitimately presupposes a 'choice negation' rather than the 'exclusion negation' that has been described in these last claims that the Gap theorist believes.

Let us turn, however, away from this discussion of the 'logic of the argument', as seen by the gap-theorist.[7] We think there is a deeper problem with this argument, also involving step (4). Let us start with the data that BOVW use to support statement (3) of their argument,

(3) n is tall $\equiv$ 'n is tall' is true.

The data reported in our Table 1 can be rearranged. Trials 1 and 2 had subjects responding to the metalinguistic statement 'what is the minimum height that you would say made "x is tall" true?' Trials 4 and 5 asked the corresponding object-language question 'what is the minimum height that x had to have in order to be tall?'.

'x is tall' is true	x is tall
178.93 cm	179.88 cm

Table 2: Metalinguistic and object language judgments for 'x is tall' (from BOVW).

Similarly, some of the data in our Table 1 can be rearranged to describe the metalinguistic and object language versions of negation, which is relevant to BOVW's

(4) n is not tall $\equiv$ 'n is tall' is false.

In Trials 1 and 2, the falsity-judgers were describing the maximum height that would lead them to say ' "x is tall" is false', and Trials 4 and 5 the falsity-judgers were describing the maximum height that would make them say 'x is not tall'. So in these cases we can compare the metalinguistic 'is false' predication with the object-language negation.

'x is tall' is false	x is not tall
165.68 cm	161.94 cm

Table 3: Metalinguistic and object language judgments for negations of 'x is tall' (from BOVW).

We see in Table 2 that the average of the positive metalinguistic claim is 178.93 cm while the average for the object language judgment is 179.88 cm. And this 0.95 cm difference is not significant, thereby providing the empirical support for (3) of BOVW's argument. But the support for (4), which is crucial to establishing the desired conclusion, is not so convincing. The difference between the metalinguistic 'false' claim and the object language negation is 3.74 cm, which (as a

[7]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 (2003).

quick check of the metalinguistic values in our Tables 2 and 3 show) is almost 30% of what subjects 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, and hence that BOVW's argument fails because of the false premise (4).

So we find that BOVW have not adequately supported the crucial claim (4), and we wish to test this claim (and other similar claims) more directly. Our results will show a place for a gap theory in the semantics of vague predicates, but that it also needs to be augmented with a 'pragmatic' account that integrates aspects of a glut theory.

3 An Experiment

The survey used for this study consisted of 20 True/False questions. The participants were presented with a synthesized image of 5 suspects in what looks like a police line-up (see Figure 3). The suspects appear to be 5′4″, 5′11″, 6′6″, 5′7″, and 6′2″, and are shown in the picture in that order[8]. The suspects were given the numbers (1-5) as names, which were printed on their faces in the image. These numbers were used to refer to the suspects in the questionnaire.

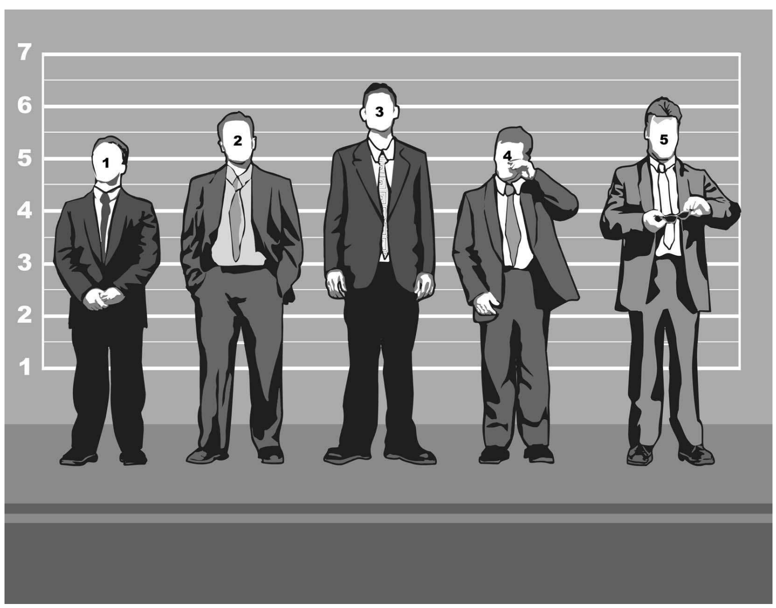

Figure 3: Suspects of Different Heights in Police Lineup

Once the participants were shown the picture, the sheet containing the 20 questions was handed out in hard-copy. The checkboxes next to each question were labeled 'True', 'False', and 'Can't

[8]The suspects were purposely not sorted by height. There were no other restrictions on their order aside from that. Both the metric measurement system and the imperial system are in common usage in western Canada.

tell'. For each suspect, (#1 for example), there were 4 corresponding questions. (With 4 questions per suspect $\times$ 5 suspects, there are 20 questions. There were no filler questions.)

#1 is tall.

#1 is not tall.

#1 is tall and not tall.

#1 is neither tall nor not tall.

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. A total of 76 subjects participated.

The data collection was done on the Simon Fraser University campus, and all participants were undergraduate Simon Fraser University students. 63.2% were native speakers of English. In total, 77.6% classified themselves as 'fluent' English speakers (which includes the native speakers), 13.2% as 'advanced', 6.6% as 'intermediate', and 2.6% did not indicate their fluency level.

Our rebuttal to BOVW draws particularly on the responses to the first two statements. Later we consider the other two sentences, in the course of presenting our own position. In Figure 4, 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 $6'6''$.

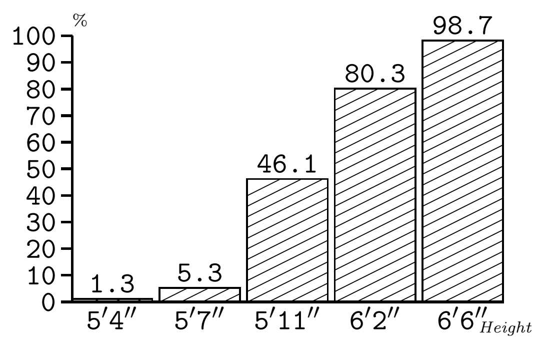

Figure 4: 'True' responses to 'X is tall'

Conversely, the percentage of *false* responses, seen in Figure 5, 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%.

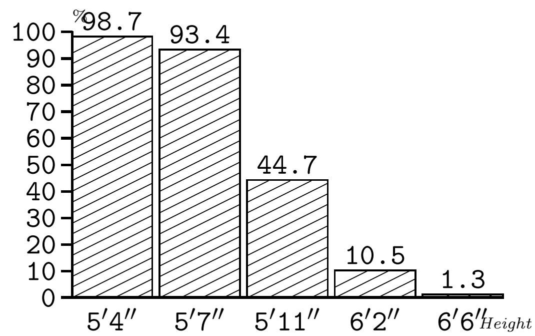

Figure 5: ‘False’ responses to ‘X is tall’

Figure 6 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″.

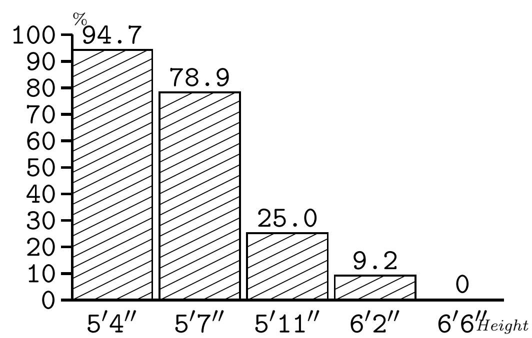

Figure 6: ‘True’ responses to ‘X is not tall’

The percentage of *false* responses to *X is not tall* is shown in Figure 7: 3.9% at 5′4″, a median of 67.1% at 5′11″, and a maximum of 100.0% at 6′6″.

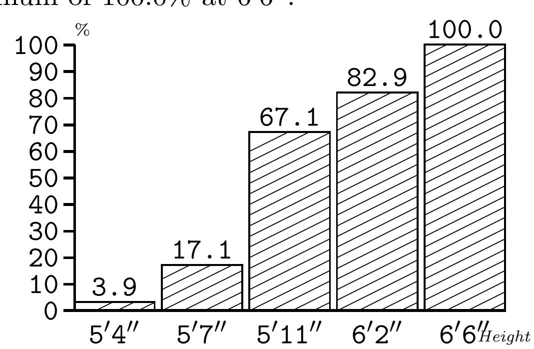

Figure 7: ‘False’ responses to ‘X is not tall’

It is the difference between these sets of answers that are problematic for the BOVW account. The numbers show a significant preference for denying a proposition over asserting its negation.[9] In classical logic, the statement ‘*a* is tall’ is true just in case its negation, ‘*a* is not tall’, is not

[9] 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%: $\chi^2(2) = 8.22; p < 0.05$.

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 untrue 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 untrue responses to 'a is not tall' than true responses to 'a is tall'.

These predictions are borne out, as shown in Figures 4–7. 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 chose to deny that #2 is tall (or not tall) when compared to the sample of those who agreed with the classical negation of each statement.

4 Discussion

We can think of three possible objections to our interpretation of the data used in our rebuttal of BOVW. The first two do not seem to be very persuasive, but the third tells against *all* the semantic accounts under consideration and shows that some 'pragmatic' effect that involves subvaluatons must be considered.

The first objection is that our use of the word 'denial' may be seen as misplaced, since the available answer in the questionnaire was 'false' and not 'wrong', 'deny', or 'disagree', etc. Suppose that a participant was in disagreement with a statement and the only three options (as in this questionnaire) were 'true', 'false', and 'can't tell', then it is obvious that the participant would have to check 'false'. So, if it is legitimate to consider 'false' as a sign of denial in this case, and we think it is, then the truth-gap approach is supported because it can be denied that an individual of borderline height is tall (or not tall) without asserting that the individual is not tall (or tall); in a gap, neither a predicate nor its negation holds, but both can be denied.

The second objection is that this could just as easily be taken as support for the 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 preferred to withhold judgement regarding uncertain cases than incorrectly attribute the predicate to them. The objector may ask why we can't say the same about our findings: if the subjects would rather deny a statement than assent to its negation, doesn't that lend support to the same assumption? The answer is yes, insofar as the participant was not presented with any other way of answering. In response, we point out that our subjects *were* given the option of checking 'can't tell'; however, 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.

The epistemicist may say that this 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 needs, then we would expect a much larger number of 'can't tell's, since this seems to be the least committing answer with regards to borderline (or uncertain) cases.

The third objection returns us to the use of negation in this experiment and in all the semantic accounts. 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 'a is not tall' is actually ambiguous: on one reading, the negation is identified with choice negation (also known as 'strong' negation and 'predicate-term' negation), in which case the statement holds if it 'super-holds', and on the other reading 'not' is identified with exclusion negation ('weak' negation, 'sentence negation'), in which case the statement holds just in case 'a is tall' does not super-hold. The objection is this: which of these two types do we think arises when we present our participants with the statement 'X is not tall'? Surely, the objector would say, if the negation was interpreted as weak negation, then there should not be a significant difference between agreeing with the statement '#2 is not tall' and denying 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 must explain why this is so.

We think the reason is pragmatic, but before we explain how pragmatics fit into the picture, we invite 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. When Lucy asks what Bill looks like, Mary provides a few answers, one of which being 'he's not tall'. John objects to the way Mary described his friend's physical stature, and in his defense he says, 'Well, he's not *not tall*. He's average.'

What this example is meant to illustrate is that the use of vague expressions, like 'tall' and 'not tall', can reasonably be assumed to accord with the Gricean principles of conversation, particularly, in this case, the maxim of quality. If it is assumed by the interlocutor(s) – our experiment subjects in this case – that this principle is observed, then it is expected that by '(not) tall' we are understood to intend the most informative reading possible, which to the hearer must correspond to that definition of '(not) tall' that s/he thinks all (or most) people would agree upon and, also, that s/he assumes that I, the speaker, 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. 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 findings that suggest a more inclusive generalization, namely, that Grice-like pragmatics govern the use of vague terms regardless of whether or not they contain negation.

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 is 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.

5.1 Data

Figures (8-11) 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. 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.

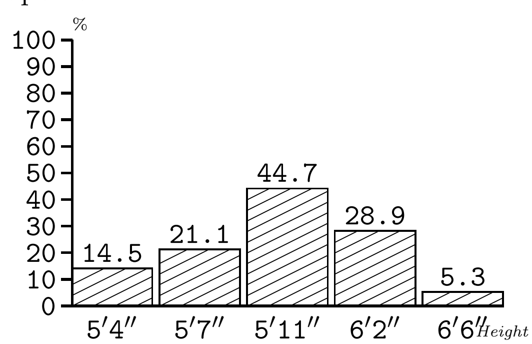

Figure 8: 'True' responses to 'X is tall and not tall'

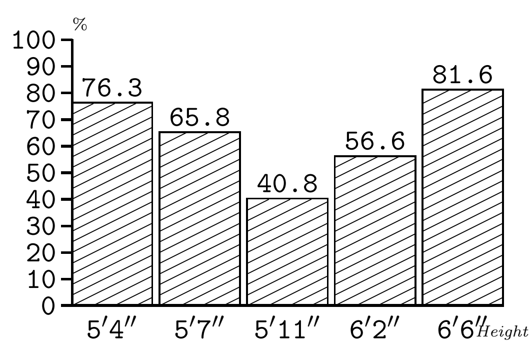

Figure 9: 'False' responses to 'X is tall and not tall'

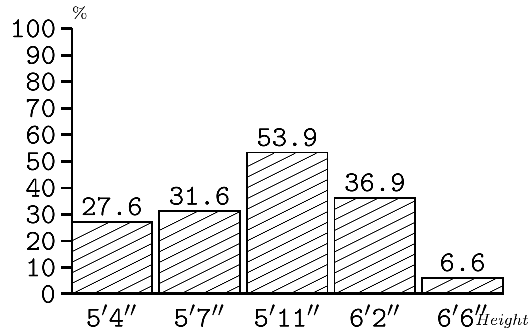

Figure 10: 'True' responses to 'X is neither tall nor not tall'

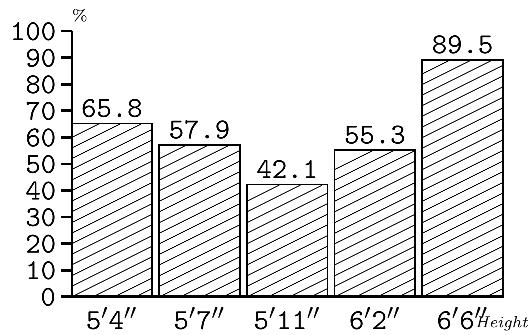

Figure 11: 'False' responses to 'X is neither tall nor not tall'

Particularly interesting, however, is how the two statements, *both* and *neither*, correlate with one another. The correlation is shown in Table 4, which shows the distribution of 'false', 'can't tell', and 'true' responses to 'x is tall and not tall' when a *true* response is given to 'x is neither tall nor not tall'. What we want to highlight is that *neither*, whose truth can justify a truth-value gap, coincides in many cases of borderline-height with *both*, which, when true, suggests a truth-value glut. For Suspect #2 ($5'11''$), for example, 53.7% of those who judged it true that he was neither tall nor not tall also thought it was true that he was *both* tall and not tall. Table 5 show the reverse correlation, namely, the distribution of truth for *neither* when *both* is thought to be true: 64.7% of those who thought $5'11''$ was tall and not tall also thought that he was neither.

neither	*both*	% within
T	F	31.7%
T	C	14.6%
T	T	53.7%

Table 4: Distribution of *both* when *neither* is true. Height = $5'11''$

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 12 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 13 illustrates the correlation in the other direction; it shows

both	*neither*	% within
T	F	35.5%
T	C	0.0%
T	T	64.7%

Table 5: Distribution of *neither* when *both* is true. Height = $5'11''$

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 $6'2''$.[10]

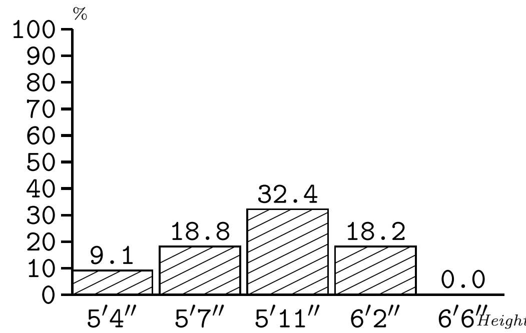

Figure 12: The Falsity of 'tall' and 'not tall' when *both* is true

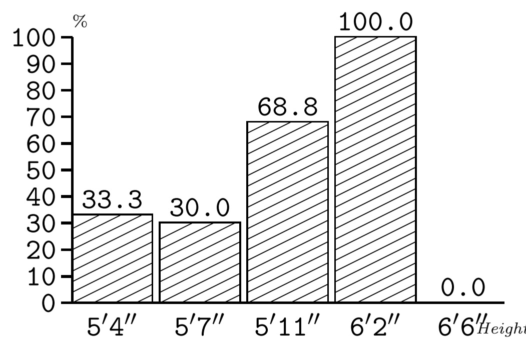

Figure 13: The Truth of *both* when 'tall' and 'not tall' are false

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

[10] It may come as a surprise that a higher percentage of this pattern was detected for $6'2''$, but the difference could originate in the overall number of subjects: the 68.8% in the case of $5'11''$ represents a total of 16 people, while the 100.0% comprises only 4.

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 of conversation, provides a way of accounting for the seemingly inconsistent patterns outlined above.

Take a simple statement like 'a is (not) tall'. Of the two interpretations, the super- and the sub-, the maxim of quantity demands that the stronger of the two be intended. If a is of borderline height, the statement is likely to be disagreed with, since a does not qualify as super-tall, or super-not-tall. Now consider a complex statement like 'a is tall and not tall'. If either 'tall' or 'not tall' is interpreted strongly, the logical result will be an empty set, since no individual can belong to the extension of 'tall' in every precisification and also belong to the extension of 'not tall' in others. Therefore, it is only the *sub*-interpretation that can make the expression meaningful. Now, in order for an individual a to be sub-tall and sub-not-tall, a would have to belong to the extension of 'tall' in some precisifications, and to its anti-extension in some other precisifications. In other words, a has to be borderline, and indeed it is mostly for suspects of borderline height that we observe an abundance of true responses to the contradictory statement.

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 $1 - [\![p]\!]$. Thus, for example, if $[\![p]\!] = 0.6$, 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 Figures 8 and 9 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 perihpery, and gradually climbs to half-truth in borderline cases. The same could be said to hold with respect Figures 10 and 11, if the disjunction of p and q is computed as $max(p, q)$. A defender of this view may add that the patterns in Figures 4-7 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.[11]

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 was 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 said to be false of a borderline individual, but in which their conjunction is said to be true. Of course, it goes without saying that further experimentation is needed before this finding can be substantialized.

(Alxatib): Linguistics, MIT
(Pelletier): Philosophy, Univ. Alberta

[11]For further criticism of the fuzzy account of vague predicates from an experimental point of view, see Ripley (2008). 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.

References

Barker, Chris. 2002. The dynamics of vagueness. *Linguistics and Philosophy* 25:1–36.

Bennett, Brandon. 1998. Modal semantics for knowledge bases dealing with vague concepts. In *Principles of Knowledge Representation and Reasoning: Proceedings of the 6th International Conference (KR-98)*, pages 234–244. San Francisco: Morgan Kaufman.

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

Cohen, Ariel, Michael Kaminski, and Johann Makowsky. 2008. Notions of sameness by default and their application to anaphora, vagueness, and uncertain reasoning. *Journal of Logic, Language and Information* 17:285–306.

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

van Fraassen, Bas. 1966. Singular terms, truth-value gaps, and free logic. *Journal of Philosophy* 63:481–495.

Frazier, Lyn, Charles Clifton, and Britta Stolterfoht. 2008. Scale structure: Processing minimum standard and maximum standard scalar adjectives. *Cognition* 106:299–324.

Graff, Delia. 2000. Shifting sands: An interest-relative theory of vagueness. *Philosophical Topics* 28:45–81.

Halpern, Joe. 2008. Intransitivity and vagueness. *The Review of Symbolic Logic* 1:530–547.

Hampton, James. 2007. Typicality, graded membership, and vagueness. *Cognitive Science* 31:355–384.

Hampton, James, Daniele Dubois, and Wenchi Yeh. 2006. Effects of classification context on categorization in natural categories. *Memory and Cognition* 34:1431–1443.

Hyde, Dominic. 1997. From heaps and gaps to heaps of gluts. *Mind* 106:641–660.

Jaśkowski, Stanisław. 1948. A propositional calculus for inconsistent deductive systems. *Studia Societatis Scientiarum Torunensis* 1:55–77. Article published in Polish with the title "Rachunek zdań dla systemów dedukcyjnych sprzecznych." English translation in *Studia Logica*, vol. 24, pp. 143–157, 1969.

Kamp, Hans. 1975. Two theories about adjectives. In E. Keenan, ed., *Formal Semantics of Natural Language*. Cambridge: Cambridge University Press.

Keefe, Rosanna. 2000. *Theories of Vagueness*. Cambridge: Cambridge University Press.

Kennedy, Chris. 2007. Vagueness and grammar: The semantics of relative and absolute gradable predicates. *Linguistics and Philosophy* 30.

Mehlberg, Henry. 1958. *The Reach of Science*. Toronto: University of Toronto Press.

Pelletier, Francis Jeffry. 1984. The not-so-strange modal logic of indeterminacy. *Logique et Analyse* 27:415–422.

Pelletier, Francis Jeffry and Robert Stainton. 2003. On 'the denial of bivalence is absurd'. *Australasian Journal of Philosophy* 81:369–382.

Priest, Graham. 2006. *In Contradiction: A Study of the Transconsistent*. The Hague: Nijhoff, 2nd edn. Expanded version of the 1987 first edition.

Ripley, David. 2008. Contradictions at the borders. Paper presented at the 2008 conference on Philosophy and Psychology of Vagueness at the Institut Jean-Nicod.

Ritov, Ilana and Jonathan Baron. 1990. Reluctance to vaccinate: Omission bias and ambiguity. *Journal of Behavioral Decision Making* 3:263–277.

Scharp, Kevin. 2005. Scorekeeping in a defective language game. *Pragmatics and Cognition* 13:203–226.

Shapiro, Stewart. 2006. *Vagueness in Context*. Oxford: Oxford UP.

Smith, Nicholas. 2009. *Vagueness and Degrees of Truth*. Oxford: Oxford UP.

Sorensen, Roy. 1988. *Blindspots*. Oxford: Clarendon Press.

Sorensen, Roy. 2002. *Vagueness and Contradiction*. New York: Oxford UP.

Spranca, Mark, Elisa Minsk, and Jonathan Baron. 1991. Omission and commission in judgment and choice. *Journal of Experimental Social Psychology* 27:76–105.

Stanley, Jason. 2003. Context, interest-relativity, and the sorites. *Analysis* 63:269–280.

Tye, Michael. 1994. Sorites paradoxes and the semantics of vagueness. *Philosophical Perspectives* 8:189–206.

Varzi, Achille. 2007. Supervaluationism and its logics. *Mind* 116:663–676.

Walton, Douglas. 2006. Using conversation policies to solve problems of ambiguity in argumentation and artificial intelligence. *Pragmatics and Cognition* 14:3–36.

Williamson, Timothy. 1994. *Vagueness*. Oxford: Routledge.

Zadeh, Lofti. 1975. Fuzzy logics and approximate reasoning. *Synthèse* 30:407–428.