The Psychology of Vagueness: Borderline Cases and Contradictions

SAM ALXATIB AND FRANCIS JEFFRY PELLETIER

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.

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 employ 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 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. After we discuss these accounts, we move to some thoughts about the methodological issues that arise in the attempt to investigate experimentally the topic of vague sentences and their negations (§2), and then to BOVW's experiment (§3).

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

---

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

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), other finitely-many valued logics (Smith, 2009), 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 вovw and will not be directly discussed here, except in passing. We describe the three approaches that we and вovw 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 multiple 'sharpenings' (also called precisifications), each of which contains some precise cut-off point. Within each sharpening, the individuals whose membership in the vague concept 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 $P(a)$ holds in every one of these sharpenings. (In the literature, these are said to be 'super true'.)

It is easiest to see this with an example that invokes an underlying ordered scale,[2] 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 180cm, another marking it at 182cm, and another at 177cm. If person a is 185cm, 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 170cm. In each precisification, b belongs to the anti-extension of *tall*, which makes the statement 'b is tall' superfalse, and therefore false. This also

---

[2] It is not so clear that every vague term can be characterized this way. A term like 'religion', for example, is vague in the sense that there are many social movements that are borderline cases of a religion. It seems that, as illustrated by the wide variety of these movements, there are many different dimensions along which a belief system might vary—the term 'religion' thus seems 'multidimensional'—and with these competing aspects it is unclear whether there is any way to compare across the dimensions. But in this paper we will stick with the sorts of vague terms that вovw consider: ones where there is a single, well-ordered dimension along which the entities being considered can vary.

makes the statement '*b* is not tall' supertrue, thus true. Now, if person *c* is 179cm, *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 the inductive part of the sorites, assumption (2) above. 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 correct 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 $P(m) \& \neg P(m-1)$, 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 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 182cm, again at 180cm, and again at 177cm. For person *a*, whose height is 185cm, 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 170cm 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 $(P(m) \ \& \ \neg P(m + 1))$, since this holds in every precisification; and hence it is True-And-Not-False that there is such an m. 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, 2001) 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 versus looking at one with $m - 1$ grains. One can't *know* that the pile with $m - 1$ grains is a heap, since it is in fact *not* a heap, according to their theory. But a pile with m grains manifests the very same evidence to the viewer, since the difference between m and $m - 1$ is indiscernible, by hypothesis; 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.[3] However, such three-valued logics require the theorist to make some seemingly arbitrary choices and various 3-valued logics 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

---

[3] We intend this to include the choice of *denying* that some sentences have a truth value at all. From the point of view of a logical system, this distinction between having a third value and not having any value at all seems to make no formal difference.

A and B we can infer "if C then D"'.[4] 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.[5] And that also violates common usage.

Since classical first order logic is the best-understood representational system, a different natural reaction sees theorists gravitating toward classical logic when they are trying to capture novel features of language—even in the case of vagueness. 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.

Of the theories we have been discussing, it is the epistemic theory that most closely adopts the classical logic 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 chide 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 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.

It should be noted, however, that having the same set of logical truths is not the same as 'being the same logic', for logic also includes the notion of a valid argument, and while supervaluation theory agrees with classical logic on the consequence relation (the same arguments are valid/invalid), it does not contain the same set of natural deduction rules as classical logic.[6] So although supervaluation

---

[4] 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.

[5] 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 (in many three-valued systems), and hence 'If p then not-p' would be true too.

[6] In particular, natural deduction rules that invoke a subproof will give different results. For, the thing that can be supposed in this subproof might not have a classical value—that is, it might be one of the supervaluationist's gappy propositions. So, assuming it 'for purposes of argument' as is done in those natural deduction argument schemas that allow this (such as conditional proof, reductio, and or-elimination) can lead to different results. A manifestation of this problem is that gap theories do not have a deduction theorem, so an argument could be valid without the corresponding conditional being unconditionally true. As Williamson (1994, p. 151) remarks, if the language contained a 'D' operator (for 'definitely') then $p \models D(p)$ would be valid ...

theory maintains the theorems and consequences of classical logic, they differ in the fine structure of arguments.

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 supervaluations, and so it is hard to see that there is any strong argument that could justifiably show that these theories are *deductively* inferior to classical logic.

2. Issues in Testing Multi–Valued Semantics and Negation

If one were to assume that there are only two truth values, True and False, and that every (declarative) sentence took one of these two values, then not only would the underlying semantics of the language be pretty straightforward, but it would be easy to discover what truth value a person thinks describes any one of these sentences. For any sentence p, one could ask whether the person thought that p was true, p was false, or whether they could not tell (or that they agree/disagree/have no opinion). And if they say that p is true, then we can legitimately infer that they think not-p is false; if they think p is false, then we can infer that they think not-p is true; and if they can't tell about p, then we infer that they can't tell about not-p either. And the reverse: if they think not-p is true/false/can't tell, then we legitimately infer that they think that p is false/true/can't tell.

2.1 Multi–Valued Semantics
But when we presume that there is some other possible semantic value — 'indeterminate', 'otherwise', 'one-half', or even 'has no truth-value at all' (='is gappy') — then matters become much more murky. Unless experimenters educate their subjects in the ways of other-valued logic in advance, how can they test whether the subjects think p is true, false, other-valued, or whether they can't tell? If the experimenter tells subjects about the third value, for example by explicitly giving them the opportunity to say 'true', 'false', 'other-valued', 'can't tell', then it is at least a possibility that they will then understand these four choices to be 'completely true', 'completely false', 'somewhere in between', and 'can't choose among the foregoing three'. And that seems to presuppose some kind of 'fuzzy theoretic'

---

that is, if p is globally true, then $D(p)$ is also. But $\models (p \supset D(p))$ isn't globally true. Recall that in supervaluation theory's analysis of the sorites, the statement $\exists n(H(n) \& \neg H(n-1))$ is super-true. Yet there is no specific value m such that $H(m) \& \neg H(m-1)$. So, existential quantifier elimination can't work either. See the discussions in Keefe, 2000 and Hyde, 1997 for the application to supervaluation/subvaluation theories of vagueness. The point was originally made in the context of supervaluation theories of presupposition in van Fraassen, 1969.

conception, which is not what was intended by introducing the single new possible answer. (Subjects might, for example, decide that only for a super-tall giant was it 'completely true' that he was tall, and only for a super-short midget was it 'completely false'. Tall people, perhaps even quite tall people, might be classified as 'somewhere in between'.) Or, they may understand the choices to mean 'I know for sure it's true', 'I know for sure it's false', 'I am not sure which one it is (maybe someone else could)', 'No one can tell about this sort of thing'. This seems to impose an 'epistemic' interpretation, which is again not what was intended by the introduction of another possible option in the answers.

The moral seems to be that *any* choice of the possible answers will itself give rise to further problems of interpretation. Of course, there is the ambiguity in the original simple choice of 'true', 'false', 'can't tell'. The third option may be understood as 'takes a third truth value' or 'I can't tell which of the first two it really is' (which presupposes a two-valued semantics and an epistemic interpretation of vagueness). And there is always the problem that some subjects may understand the 'true' answer to mean 'true and I know it' while others might understand it as 'true and anyone would agree'. Any of these possible understandings could characterize some part of the subject population, while other parts are better characterized by different ones of the understandings.

2.2 Negation

Matters are further complicated by the fact that there are at least three different conceptions among logicians about how negation should work in three-valued logics.

Option one, which we will symbolize $\sim p$, is choice negation (also called 'strong negation' and Łukasiewicz negation, meaning 'I take the opposite truth value from p', where only the classical True and False have opposites). Choice two, which we symbolize $\neg p$, is exclusion negation (also called 'weak negation', meaning 'p has some value other than true'); and choice three, which we symbolize $-p$, is intuitionistic negation (also called 'Gödel negation', meaning 'p has the value false'). Using 'O' as the name of the third, 'other', truth value (and including the intuition that it could mean 'has no truth value at all'), the three negations are described in Table 1.

In the presence of gaps, supervaluation theory predicts that the negation of a gappy sentence is also a gap, since a gap means 'true in some but not other precisifications' while its negation means 'is false in the former and true in the latter precisifications'. Supervaluation theory therefore predicts that people are employing choice negation. And so they might give 'true, false, neither true nor false, don't know' as alternatives when they construct an experiment to question subjects about what truth value some sentence has.

But if we wanted to remain neutral about the existence of this third value, hoping perhaps to *discover* whether people believed there was a third value, and also allowing for the possibility that they just don't know what value some proposition has, then

p	$\sim p$	$\neg p$	$- p$
T	F	F	F
O	O	T	F
F	T	T	T

Table 1 *Truth-tables for choice, exclusion, and intuitionistic negations.*

using tests about 'true', 'false' and 'don't know' together with negative statements is bound to lead to difficulties. If we wish to remain neutral, then we can't really give them the information that there *is* another truth value to consider. But if we don't do that, then we cannot distinguish among any of our three negations.

Note that it is only supervaluation theory that thinks that a negated sentence can take a third semantic value; both exclusion negation and intuitionistic negation assign one of the classical values to all negations. Serchuk *et al.* (2008) tested whether the difference in English between 'not $F(a)$' (e.g. 'Tim is not tall') and 'it is not the case that $F(a)$' (e.g. 'It is not the case that Tim is tall') would distinguish choice from exclusion negation: they thought that when given a borderline case of $F(a)$ (thus having value O), they could test which negation was used by asking subjects about their views of sentences like 'not $F(a)$' versus 'it is not the case that $F(a)$'. They report that subjects answered 'something other than ''true'' or ''false''', for the former, and answered 'true' for the latter. It seems to us quite difficult to determine from the fact that they used 'something other than ''true'' and ''false''' whether they were really using choice negation, or were perhaps adverting to not being able to tell. Furthermore, the tests that involved more complex sentences such as 'Either $F(a)$ or not $F(a)$' versus 'Either $F(a)$ or it is not the case that $F(a)$' (and others) gave results that were not consistent with *any* of the possible negations, thereby casting doubt as to whether the assumed interpretations of 'not $F(a)$' versus 'it is not the case that $F(a)$' were really adopted and used by their subjects (as Serchuk *et al.* acknowledge).

Interfering even further with the attempt to describe subjects' use of negation as one or the other of our three negations is the fact that there could be further, pragmatic processes that cause the subjects to answer in whatever way they do. It could be that the *semantic* meaning of negation is given by one or another of the three negations, but this feeds into some pragmatic (non-semantic) rationale for giving a final answer to any experimental question. And this final answer might be the same as what a different one of the semantic meanings for negation predicts. For example, we note that there is an alternative, pragmatic explanation of the results that Serchuk *et al.* do have (which they assume to be semantic facts). We might say that the Gricean maxim of quantity predicts that any expression containing a vague predicate ought to be assigned 'the strongest interpretation' by a hearer. So, this maxim predicts that sentences of the form 'a is not F' will be interpreted as 'a is choice-not F' by hearers. This account agrees with the Serchuk *et al.* view in its

predictions about '*a* is not *F*', but it appears to differ with sentences of the form 'it is not the case that *a* is *F*'. For, if, as they say, this really expressed exclusion negation, then it could not have been assigned the strongest possible interpretation, for in that case it too should be have been interpreted as saying '*a* is choice-not *F*'. But this is apparently not what Serchuk *et al*. found. The imagined pragmatic account could say that there is here a conflict between the maxim of quantity and the maxim of manner. For, given that '*a* is not *F*' is what expresses '*a* is super not-*F*', then a speaker would not use 'it is not the case that *a* is *F*' to express that proposition. So something other than the super-false interpretation must be intended. And some sort of accommodation might be called upon to signal that it is a different form of negation, not to be treated as the classical negation in a gappy environment. Whatever the details, the point is that it can be impossible to determine whether a subject is answering in accordance with the semantic meaning of a negation or with a pragmatic extension of some different semantic meaning.

2.3 Our Methodology

In our studies below we take a different approach to this vexed topic, but allow that it deserves further study. The way we have chosen to investigate this issue is to compare the truth (as frequency of truth-predications) of a proposition using a vague term, and the falsity (as frequency of falsity-predications) of its negation. If tallness, for example, were a sharp concept, and subjects were presented with a borderline case of vagueness, one would expect an equal number of truth-judgers of '*a* is tall' and falsity-judgers of '*a* is not tall' (together, perhaps, with a number of 'can't tell's). If, however, we find significantly more falsity-judgers of '*a* is tall' than truth-judgers of '*a* is not tall', this suggests that more people are willing to reject—by judging it false—the borderline proposition '*a* is tall' than agree with—by judging it true—the proposition '*a* is not tall', and we should therefore say that the falsity judgment is tracking intuitionistic negation. Likewise, one would also expect an equal number of truth-judgers of '*a* is not tall' as falsity-judgers of '*a* is tall' in the classical case. But if we find significantly more falsity-judgers for '*a* is not tall' than truth judgers for '*a* is tall', then again we should say that the falsity judgment is tracking intuitionistic negation.

As we will see below, this is in fact what we find. We admit, however, that these results do not square exactly with those of Serchuk *et al*. But we think our account covers as much of the data as does the supposition made in Serchuk *et al*., and in addition can account for some puzzling data in ʙovw that were not remarked upon in their discussion. We postpone this until our evaluation section (§3.3).

We draw the reader's attention to the fact that we are leaving it open for a theory that incorporates more than one kind of negation. Consideration of sentences like '*a* is not *not* tall' give some added reasons to explore this option, and we will do so in §§5, 6.2.

3. BOVW's Experiment

3.1 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 employment* (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.

	Trial 1	Trial 2	Trial 3	Trial 4
Truth-judgers	178.30cm	179.55cm	181.49cm	170.28cm
Falsity-judgers	167.22cm	164.13cm	160.48cm	163.40cm

Table 2 *Truth and falsity judgments for 'x is tall' (from BOVW).*

3.2 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 2.[7]

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 '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.

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 *x*s 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

---

[7] 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.
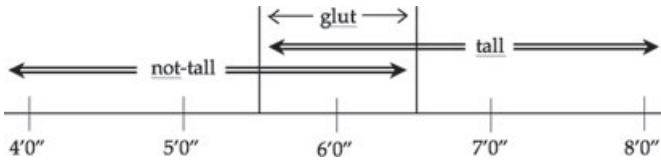
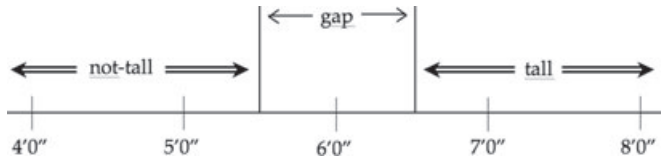**Figure 1** *Subvaluation predictions about height*

Figure 2 *Supervaluation predictions about height*

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.

3.3 Evaluation

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.

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.[8] 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.

---

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

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.

3.3.1 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.

The treatment of higher-order vagueness varies across theories, but in superval-uations 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 200cm is certainly not admissible, and the corresponding precisification is therefore not considered when assessing the tallness of a given individual.[9] Similarly, 160cm 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 200cm and 180cm, 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

---

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

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.

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' (p. 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.

3.3.2 The Absurdity of Denying Bivalence.

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 178.30cm, and in the second (Study 2), the average was 179.55cm. In comparison, the average answer of truth-judgers in the first *non*-metalinguistic trial was 181.49cm (Study 4), and 178.28cm 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 167.22cm, and in the second trial (Study 2) the average was 164.13cm. In Studies 4 and 5 (non-metalinguistic), the average answers for falsity-judgers were 160.48cm and 163.40cm, 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 (pp. 388−389) to show that this cannot be correct (they wish the '=' 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 = '*n* is tall' is true (as shown by their experimental results)
    (4) *n* is not tall = '*n* is tall' is false (as shown by their experimental results)
    (5) *n* is not tall = '*n* is tall' is not true (from equivalence (3))
    (6) *n* is not not tall = '*n* is tall' is not false (from equivalence (4))
    (7) *n* is tall = '*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 therefore 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 gap theorists generally, and supervaluationists in particular, will find this general form of argumentation unsatisfactory, and will object to the reasoning involved in a number of the individual steps of the argument. 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, object-language 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 steps (3) and (4) 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. Premise (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 supervaluationist thinks that 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 we above (§2.2) called 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.

Not only do we think that supervaluationists believe there to be flaws in the argumentation, but we find that *any* theorist who believes there to be a third value besides True and False will object to it. Their objections, however, will be different depending on the sort of negation they think there to be. So, let us consider more generally this line of response made by believers in three semantic values. We note first that the *metalinguistic negation* of a sentence's having a certain truth value—for instance, when it is employed to say 'p is not true'—*must* be understood as exclusion negation, saying that p takes one of the truth values other than True. In our imagined three-valued case, such a claim amounts to saying 'p is either False or Gap'. To say 'p is not False' is similarly to say 'p is either True or Gap'. This means that the negations on the right hand sides of (5)−(7) are exclusion negations. To say otherwise is simply to insist, question-beggingly, on a two-valued system.

Any of our different three-valued logicians will agree with what they take BOVW's premise (3) to be saying, namely the conjunction of these two claims:

(a) $p \models \textit{True}(p)$
(b) $\textit{True}(p) \models p$

But now consider the following, which might be thought to follow from (a) and (b) respectively by a kind of 'argument modus tollens':

(c) $\textit{not-True}(p) \models \textit{not-}p$

   (d) *not-p* $\models$ *not-True(p)*

Since in the three-valued case, 'not-True(p)' is a kind of shorthand for 'either False(p) or Gap(p)', these two claims are equivalent to:

   (c') *(False(p) or Gap(p))* $\models$ *not-p*
   (d') *not-p* $\models$ *((False(p) or Gap(p))*

Our different views of negation will think of (c') and (d') differently. If 'not-p' is interpreted as $-p$ (that is, as intuitionistic negation), then (c') will be seen as incorrect, since when p has the value Gap (and hence the premise is true), then $-p$ has the value False. Furthermore, even if 'not-p' is interpreted as $\sim p$ (choice negation), then again (c') will be judged to be incorrect, since when p has the value Gap (and hence the premise is true), then $\sim p$ has the value Gap (which is not true). We repeat the truth table from §2.2 for assistance in verifying the claims made in this and the next paragraphs:

p	$\sim p$	$\neg p$	$-p$
T	F	F	F
G	G	T	F
F	T	T	T

Table 1 (repeated) *Truth-tables for choice, exclusion, and intuitionistic negations.*

   Since (c') is incorrect for both intuitionistic and choice negation, if the translation of object-language negation into the metalanguage is one of them, then the generation of the contradiction using it is of no concern to such theorists, since they just find that one of the premises of the argument is wrong. (Note that (c') is the right-to-left direction of line (5) of the ʙᴏᴠᴡ proof. Without (5), line (9) cannot be derived; and that is half the alleged contradiction.)[10]

   Of course, all three versions of negation will assent to (d'), since whenever not-p is true using any type of negation, then p is either False or Gap, as can be verified by the Table. But for choice negation ($\sim p$) and intuitionistic negation ($-p$) something stronger holds:

   (e) *not-p* $\models$ *False(p)*.

(e) does not hold for exclusion negation, however, since $\neg p$ can take the value True even when p takes the value Gap. Since (e) does not hold for exclusion negation, if the translation of object-language negation into the three-valued metalanguage is exclusion negation, then the left-to-right direction of (5) is wrong. And without

---

[10] We might also note that the inference from step (6) to (7)—'double negation elimination'—is not valid using intuitionistic negation. For, if 'n is tall' has value G, then 'n is not not tall' has value T. This shows that the inference from the latter to the former is invalid.

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

Table 3 *Metalinguistic and object-language judgments for 'x is tall' (from* BOVW*).*

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

Table 4 *Metalinguistic and object-language judgments for negations of 'x is tall' (from* BOVW*).*

this, once again line (9) of the BOVW proof can't be derived, and the contradiction doesn't follow.

Thus, no matter which form of negation is seen as being in the object language, a three-valued theorist will say that its translation into the metalanguage will show that one or another of the crucial premises of the BOVW argument is incorrect. It gets any apparent strength it has by surreptitiously importing a bivalent semantics into the argument. And so the argument is beside the point.

Let us turn, however, away from this discussion of the 'logic of the argument', as seen by the gap-theorist.[11] 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 = 'n is tall' is true.

The data reported in our Table 2 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?'. Table 3 puts all these figures together:

Similarly, some of the data in our Table 2 can be rearranged to describe the metalinguistic and object language versions of negation, which is relevant to BOVW's:

(4) n is not tall = '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, and we have Table 4.

---

[11] 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.

We see in Table 3 that the average of the positive metalinguistic claim is 178.93cm while the average for the object language judgment is 179.88cm. And this 0.95cm 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 3.74cm, which (as a quick check of the metalinguistic values in our Tables 2 and 4 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 with an intuitionistic negation 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.[12]

4. 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.[13] 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.

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 tell'. For each suspect, (#1 for example), there were 4 corresponding questions. (With 4 questions per suspect × 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

---

[12] We re-emphasize: supervaluation theory is only one sort of gap theory. We argued in §2.2 that they are committed to a choice ('strong') negation; but other gap theories are not committed in that way. Some could adopt an exclusion ('weak') negation, others adopt an intuitionistic ('Gödel') negation, and (as we will do) some could adopt more than one type of negation.

[13] 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.

Figure 3 *Suspects of differing heights in police lineup*

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″.

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 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″.
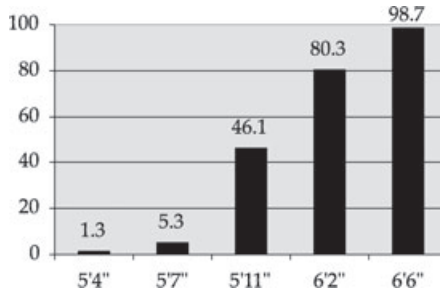
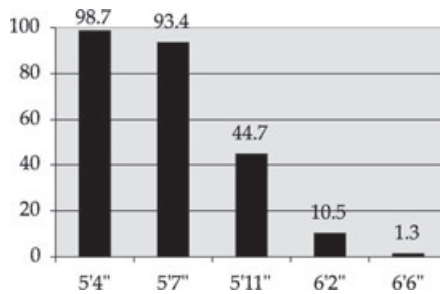**Figure 4** *'True' responses to 'x is tall'*

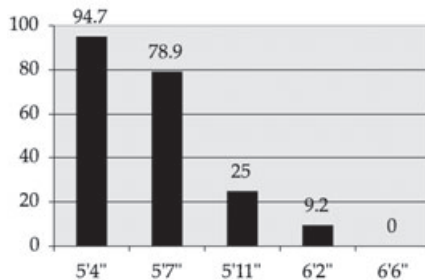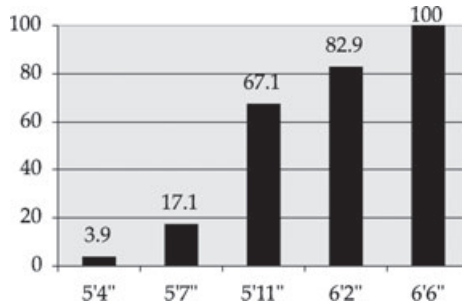**Figure 5** *'False' responses to 'x is tall'*

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 $5'1''$, and a maximum of 100.0% at $6'6''$.

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.[14] In classical logic, the statement '*a* is tall' is true just in case

---

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

Figure 7 *'False' responses to 'x is not tall'*

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 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. Similarly, 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 both that #2 is tall and that #2 is not tall, when compared to the sample of those who agreed with the classical negation of each statement.

ʙᴏᴠᴡ might claim that this could as easily be taken as support for their epistemic hypothesis. Recall that ʙᴏᴠᴡ 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. 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 a 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, 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.

A referee remarks that 'can't tell' may be interpreted as a regrettable lacuna in one's knowledge, 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 §6, 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 refer the reader to our discussion in §2.3.

5. Discussion

We can think of two further possible objections to our interpretation of the data used in our rebuttal of BOVW. The first does not seem to be very persuasive, but the second tells against *all* the semantic accounts under consideration and shows that some 'pragmatic' effect that involves subvaluations 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. In response to this objection, we first offer a hypothetical scenario in which our participants are given the answers 'true', 'not true', and 'can't tell' as options (instead of 'true', 'false', 'can't tell'). If, under this version of the experiment, we were to find the same distribution of answers and the same preference for checking 'not true' to '#2 is tall' (and '#2 is not tall') over checking 'true' to '#2 is not tall' (and '#2 is tall'), then what we claim to support gap-theories becomes more obvious. In a gap theory, 'not truth' includes both falsity and gap-hood, and this makes a statement that predicates tallness of an individual of borderline height a not-true statement. It would then be reasonable to say that a gap theory predicts the preference described above to emerge, since the theory will predict significantly more 'not true' responses to both 'x is tall' and 'x is not tall' when x's tallness is gappy (when x is of borderline height, that is). So, the statistical results would in fact show support for a theory like that of supervaluations. But now suppose that a participant was in disagreement with some statement, and was presented with only three possible answers: 'true', 'false', and 'can't tell' (as in our questionnaire). In this case, it is quite plausible to expect that the participant will check 'false'; for that is the best substitute for 'not true' among the available

answers. And once we allow for 'false' to be an indicator of denial in this case, we see how the data can be taken to support gap-theories, since on a gap-theoretic approach it can be *denied* that a borderline individual is tall without *asserting* that the same individual is not-tall. (And *mutatis mutandis* for the case of denying that the individual is not-tall.) In other words, a supervaluationist could take denial to signal not-supertruth but take explicit negation to be choice ('strong') negation: in a gap, neither a predicate nor its negation holds, but both can be denied. (Of course, this unsupported claim of ours needs to be bolstered by experimental evidence. One could re-run our experiment with the replacement of 'false' by 'not true'. We predict the same pattern of answers as we obtained.)

The second objection returns us to the use of negation in this experiment and its role 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 three-ways ambiguous: on one reading, the negation is identified with choice/strong negation (in which case '*a* is not tall' holds if it 'super-holds'), on a second reading 'not' is identified with exclusion/weak negation (in which case the statement holds just in case '*a* is tall' does *not* super-hold), and on a third reading 'not' is identified with intuitionistic negation (in which case '*a* is not tall' holds if it super-holds).[15] The objection is this: which of these three 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. Although this seems like a plausible argument, it is not a complete analysis of the case, since there are different negations that might be in play, as we will now try to show.

We think that there is a pragmatic reason behind the emergence of the choice/intuitionistic interpretation for 'not', but before we explain, we invite the reader to consider the following scenario, which we alluded to in §2.3. 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

---

[15] So, the differences between choice and intuitionistic negation would show up only in regards to multiply-negated sentences. In choice negation, '*a* is not not tall' would hold if it superholds, i.e. if '*a* is tall' superholds. With intuitionistic negation, '*a* is not not tall' holds if it superholds, but here this can happen in two ways: '*a* is tall' might be true or it might be a gap. This is why the inference from (6) to (7) in BOVW's argument of §3.3 is seen differently by choice–negation versus intuitionistic negation theorists.

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.' In the earlier section, we suggested that this shows that there must be two different 'not's involved, for otherwise John's comment would merely be equivalent to 'he's tall'.

However, the situation is somewhat more complicated than this observation might suggest on the surface. Suppose the person is of average-height, and we are considering a three-valued system (or a two-valued-plus-gap system). Then the statement 'He is tall' takes the third value. So, when Mary says 'He's not tall', if she is to be speaking truthfully[16] she must be using exclusion negation, for that is the only negation that will convert an 'other-valued' statement into a truth. (Of course, we might be happy to say that Mary is speaking in the 'otherwise' manner, rather than truthfully. In that case she must be using choice negation, since that's the only negation that would yield an 'other' for the negation of an 'other'.) When John tries to correct Mary, or correct the impression left by Mary's statement, there are two ways to understand what he is doing.[17] One way is that 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 exclusion negation) or negating an 'other' (if Mary was using choice negation). Note in these cases that if Mary were using 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 Mary could not be using exclusion negation (nor, from the previous consideration, could she be using intuitionistic negation). So maybe Mary was using choice negation and John is denying the 'other' value to Mary's claim. This would mean that John was using exclusion negation, since that is the only negation that takes an 'other' value into truth. The second way of understanding what John is doing is this: he is *not* trying to negate Mary's negation, but rather is trying to use his own negation twice to get the idea across. In such a scenario, John must be using intuitionistic negation both times, since that is the only double negation of an 'other' that will generate a true.

From this analysis we see: (from the first understanding) it must be that Mary is using choice negation and John is using exclusion negation ... which is the understanding given in the previous paragraph. But (from the second understanding): Mary could be using either intuitionistic or choice negation (and she would be speaking either truly or 'otherwisely', respectively); however, John must then be using intuitionistic negation in both occurrences, to be speaking truly. As with most attempts to analyze how ordinary (or experimental) employment of

---

[16] '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.

[17] Again, we want John also to be speaking truthfully, in the semantic sense of the term. His worry is that Mary's semantic truthfulness might be pragmatically interpreted in an inappropriate way.

multiple values and negations interact, we see here that there are different groups of choices that can direct our analyses.

Whatever logical form(s) of negation one chooses to employ in representing natural linguistic negation, it seems that one ought to complement the semantics with a pragmatic account of how negation is used in conversation. In particular, when an experimenter presents a group of participants with questions or statements that contain negated (or even unnegated) vague expressions, it can be assumed quite reasonably 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 the experimenters intend that this principle will be observed by the participants, 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. 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, an account we describe in §6.2.

6. Contradictions and Borderline Cases: Gaps versus 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.

6.1 Data

Figures 8 and 10 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, as shown in Figures 9 and 11. 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*.
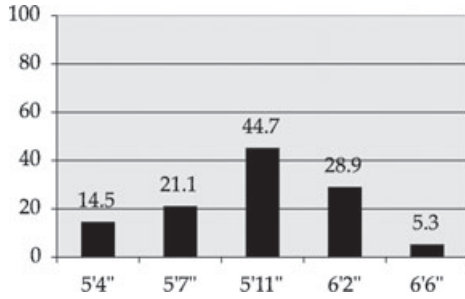
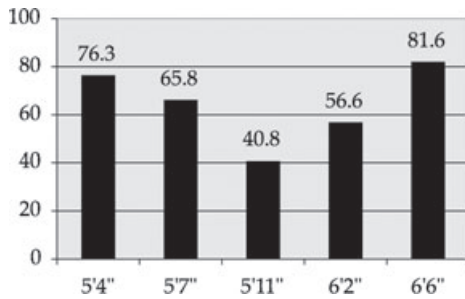**Figure 8** *'True' responses to 'x is tall and not tall'*

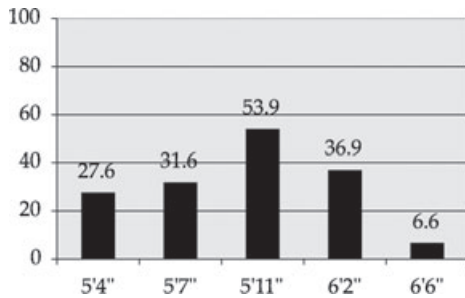**Figure 9** *'False' responses to 'x is tall and not tall'*

Figure 10 *'True' responses to 'x is neither tall nor not tall'*

Particularly interesting, however, is how the two statements, *both* and *neither*, correlate with one another. The correlation is shown in Table 5, 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 (more than half!) 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*
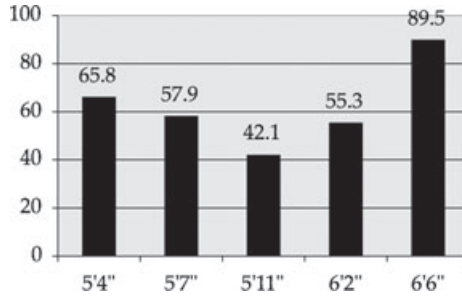
Figure 11 *'False' responses to 'x is neither tall nor not tall'*

Neither	Both	% within
T	*F*	31.7%
T	*C*	14.6%
T	*T*	53.7%

Table 5 *Distribution of* both *when* neither *is true. Height = 5′11″*

Both	Neither	% within
T	*F*	35.5%
T	*C*	0.0%
T	*T*	64.7%

Table 6 *Distribution of* neither *when* both *is true. Height = 5′11″*

tall and not tall. Table 6 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 both tall and not tall also thought that he was neither.

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 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″.[18]

---

[18] 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'.
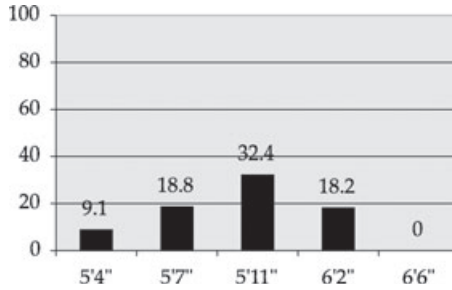
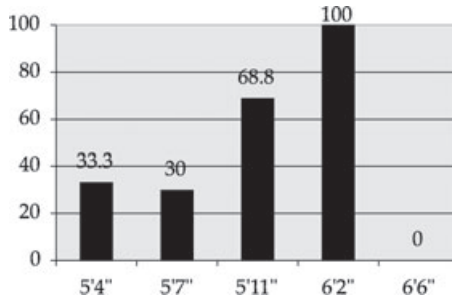**Figure 12** *The falsity of 'tall' and 'not tall' when* both *is true*

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

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

Table 7 *Percentage of Ss who gave same answer to '#2 is tall' and '#2 is not tall'.*

Other interesting findings center around our borderline suspect #2. For example, we find the difference between the number of subjects who think both '#2 is tall' and '#2 is not tall' are both true versus those who think they are both false (Table 7). And it is also interesting to note the number of subjects in each of these categories who think '#2 is neither tall nor not tall' (Table 8).

Of our 76 subjects, 34 (44.7%) thought that '#2 is tall and not tall' was true (compare this with how few thought both conjuncts were true!); 31 (40.8%) thought it false; and 11 (14.5%) answered 'can't tell'. So, more than half of our subjects (the 'true' and the 'can't tell') are not "classicalists". We might wish to look within each of these groups to see what they thought of the individual conjuncts: are they both T, both F, both C, one T and one F, or other (one with a T/F value and the other C)?

Both	Pattern on conjuncts	% within answer type of *both*
	Both *T*	2.9%
	Both *F*	32.4%
T	Both *C*	5.9%
	One *T*, one *F*	55.9%
	Other	2.9%
	Both *T*	3.2%
	Both *F*	12.9%
F	Both *C*	0.0%
	One *T*, one *F*	74.2%
	Other	9.7%
	Both *T*	9.1%
	Both *F*	9.1%
C	Both *C*	9.1%
	One *T*, one *F*	54.5%
	Other	18.2%

Table 8 *Percentage of Ss within type of answer to '#2 is tall and not tall'.*

Of the subjects who thought '#2 is tall and not tall' was true, most thought that both conjuncts assumed classical truth values (the majority of these, 58%, thought that the positive was true while 42% thought the negative was the true one). The next-most common view, 32.4%, thought that both conjuncts were false. Besides violating classicality by claiming that *both* was true, it seems that they also would not agree to the rule of *and*-elimination.

Let us turn our attention now to the *neither* question. Of our 76 subjects, 41 said that '#2 is neither tall nor not tall' was true; 32 said it was false and 3 gave the answer 'Can't tell'. (It seems therefore that 57.9% of our subjects are 'non-classical' when it comes to predications of the 'neither F nor not-F' sort—they either think it is false or they can't tell, while only 42.1% are 'classical'.) Within each of the first two of these three categories, we wish to know the percentage of subjects who thought that each disjunct was true, thought each disjunct was false, thought they couldn't tell about each one, thought one was true and the other false, and the remainder (who thought that one disjunct had a classical value but the other they couldn't tell about), which is given in Table 9.

Of our *neither*-is-true group, we see that more than half of them nonetheless think that one or the other of the disjuncts is true and the other false, which seems to violate yet another classical law or rule of inference. Meanwhile, in the "classicalist" group—those who think that *neither* is false—we see that almost 80% of them are also classical in thinking that one of the disjuncts is true and the other false. Another 15% are not quite as classical because they think both disjuncts are false, which they shouldn't do on other classicalist grounds.

Neither	Pattern on disjuncts	% within answer type of *neither*
	Both T	2.4%
	Both F	24.4%
T	Both C	9.8%
	One T, one F	51.2%
	Other	9.8%
	Both T	6.2%
	Both F	15.6%
F	Both C	0.0%
	One T, one F	78.1%
	Other	0.0%

Table 9 *Percentage of Ss within type of answer to '#2 is neither tall nor not tall'.*

6.2 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 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.

We now propose a rough-and-ready account of the data that does away with the logical complications that accompany negation, and instead treats negation as denoting an operation which, when applied to a set S, simply returns the

complement of S. It was already suggested that the pattern where 'a is tall' and 'a is not tall' are both judged false and where, simultaneously, 'a is tall and not tall' is judged true, can be accounted for by assuming that 'tall' is associated with both a super- and sub-interpretation, and, in addition, that the emergence of either interpretation is governed by familiar Gricean pragmatic factors. 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[19] is interpreted normally, as some subset of $\mathcal{D}$ (the ones that manifest the property). 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.

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.[20] 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, 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

---

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

[20] We see this as an application of the 'Strongest Meaning Hypothesis' of Dalrymple *et al.* (1998).

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 (through the denotation of 'and') two sets. The first element is the intersection of super-tall and its complement, the second is that of sub-tall and its complement, the third is that of super-tall and the complement of sub-tall, and the fourth is that of sub-tall and the complement of super-tall. It seems quite reasonable to include in the formulation of Quality, whatever that may be, a ban on interpreting predicates as (trivially) empty, and if this possibility is pursued, Quality would block every one of the listed interpretations except the last, i.e. it admits the intersection of sub-tall and the complement of super-tall, which is just the set of individuals who are sub-tall and sub-not-tall, namely, the borderline individuals.

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 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 $p \& q$ is defined as the minimum of the truth values of the conjuncts p and p. If the truth-value of p were 0.6, for example, and the value of p were 0.3, then the value of $p \& q$ will be $min(p, q)=0.3$. This makes it possible for contradictory expressions like $p \& \neg p$ 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 $p \& \neg p$ 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 periphery, 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 p 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.[21]

---

[21] 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

7. Issues for Future Work

In §2.1 we discussed some of the issues involved in trying to gather empirical evidence about subjects' truth-value judgments when we wish to leave open whether there are more than two truth values and also wish to offer the opportunity to say that they can't tell. In §2.2 we brought forth further issues involving negation and its interaction with the ways one might attempt to test for the possibility of more than two truth values. We now return to those issues, relating them to the data we have presented in the last section.

What seems really interesting in our data is that hardly anyone says 'can't tell' when asked about the tallness or non-tallness of #2. One way to interpret this is that the bare choice of 'true/false/can't tell' does not distinguish clearly enough between 'I don't know which one of the two truth-values' versus '#2 takes some other truth value'. If this is right, then it is tempting to move to a position which understands the independently-presented question we asked, '#2 is neither true nor false', as their attempt to access this third value. (So, if they say true to the *neither* question, then they are attributing the intermediate value, or a gap, to #2's tallness).

But if we do this, we then have a difficulty in understanding what their negation is doing. For, if they say T to *neither*, and we interpret this as saying 'there's a gap/third value to #2's tallness', we are faced with the question of what value we think they should then give to the '#2 is tall' and '#2 is not tall' questions. Should it be 'can't tell' on the grounds that they would really like to say 'neither true nor false' but weren't given the opportunity to do so? Our results show that such a theory is not right: of our 76 subjects, only 4 answered 'can't tell' to the individual disjuncts and also answered 'true' to the *neither* question. So, this is not the correct way to look at the subjects' internal view of their choices. But in turn, this seems to show that people do not tend toward a supervaluation-style theory, since that is the theory that is naturally associated with this view.

Of course, a supervaluationist might reply that it is wrong to say they are committed, in this sort of forced-answer situation, to having the individual disjuncts necessarily take the value 'can't tell'. After all, the theorist may say, subjects were unhappy that they could not answer 'neither' to those questions, and perhaps they just randomly picked one of the three possible answers. Or perhaps they used some other, irrelevant-to-the-supervaluation-theory method to choose an answer. In that case, our supervaluation theorist may say, we should allow *any* answer to the disjuncts. The thing of importance is that they thought that '#2 is neither tall nor not tall' is true.

Perhaps. However, it might also be noted that of the 41 subjects that answered 'true' to *neither*, 22 of them (53.7%) also answered 'true' to '#2 is both tall and

---

(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.

not tall', which should not happen, according to supervaluationism. And 6 (14.6%) of them answered 'can't tell' to this, which also shouldn't happen according to supervaluationism. Thus only 13 of the subjects that answered 'true' to *neither* (31.7%) are acting as supervaluationists, even with this generous understanding of the way they might answer the atomic questions. (And recall that this means that only 13 of 76 subjects in the entire pool, 17.1%, were answering in accordance with supervaluationism.) Of course, this is all conditional on the idea that answering '#2 is neither tall nor not tall' as true defines supervaluationism, which is a stretch.

Another interesting and possibly relevant fact is that more than 90% thought that '#2 is tall' is either true or is false. And more than 90% thought that '#2 is not tall' is either true or is false. Indeed, 63% of the subjects were 'classical' about their judgment of the two basic statements: they said one of the two was true and the other false. The next biggest group thought they were both false (16 of them: 21.1%). Only 9 of them answered with at least one 'can't tell' (12%), of which 4 had two 'can't tell' (5%).

So, we can see that the 'can't tell' answer isn't doing very much, and it certainly is hard to interpret 'can't tell' as 'takes a third value (or a gap)'. A future experiment should probe all the different ways one might try to tell whether or not the subjects are adverting to a third value (including a gap), including the possibility of alerting them to its possible existence with a possible answers like 'assumes some other truth value', 'has no truth value', or 'neither true nor false' (as a possible answer, rather than being asked to judge whether it is true or false. And as we mentioned in §5, one would also like our experiment re-run with 'not true' in place of 'false'.

Also here, we need to acknowledge 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.

8. 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. Of course, it goes without saying that further experimentation is needed before this finding can be substantiated.

Department of Philosophy and Linguistics
MIT

Department of Philosophy
University of Alberta

References

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

Bennett, B. 1998: Modal bases. In *Principles of Knowledge Representation and Reasoning: Proceedings of the 6th International Conference (KR-98)*. San Francisco: Morgan Kaufman.

Bonini, N., Osherson, D., Viale, R. and Williamson, T. 1999: On the psychology of vague predicates. *Mind & Language*, 14, 377–93.

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

Dalrymple, M., Kanazawa, M., Kim, Y., Mchombo, S. and Peters, S. 1998: Reciprocal expressions and the concept of reciprocity. *Linguistics and Philosophy*, 21, 159–210.

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

van Fraassen, B. 1966: Singular terms, truth-value gaps, and free logic. *Journal of Philosophy*, 63, 481–95.

van Fraassen, B. 1969: Presuppositions, supervaluations, and free logic. In K. Lambert (ed.), *The Logical Way of Doing Things*. New Haven, CT: Yale University Press, 67–91.

Frazier, L., Clifton, C. and Stoltetfoht, B. 2008: Scale structure: processing minimum standard and maximum standard scalar adjectives. *Cognition*, 106, 299–324.

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

Halpern, J. 2008: Intransitivity and vagueness. *The Review of Symbolic Logic*, 1, 530–47.

Hampton, J. 2007: Typicality, graded membership, and vagueness. *Cognitive Sciences*, 31, 355–84.

Hampton, J., Dubois, D. and Yeh, W. 2006: Effects of classification context on categorization in natural categories. *Memory and Cognition*, 34, 1431–43.

Hyde, D. 1997: From heaps and gaps to heaps of gluts. *Mind*, 106, 641–60.

Jaśkowski, S. 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*, 24, 143−57, 1969.

Kamp, H. 1975: Two theories about adjectives. In E. Keenan (ed.), *Formal Semantics of Natural Language*. Cambridge: Cambridge University Press: 123−55.

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

Kennedy, C. 2007: Vagueness and grammar: The semantics of relative and absolute gradable predicates. *Linguistics and Philosophy*, 30, 1−45.

Kyburg, A. 2000: When vague sentences inform: a model of assertability. *Synthèse*, 124, 175−91.

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

Pelletier, F. J. 1984: The not-so-strange modal logic of indeterminacy. *Logique et Analyse*, 27, 415−22.

Pelletier, F. J. 1988: A jó, a rossz és a csúf. *Tertium non Datur*, 5, 89−115.

Pelletier, F. J. and Stainton, R. 2003: On 'the denial of bivalence is absurd'. *Australasian Journal of Philosophy*, 81, 369−82.

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

Ripley, D. 2008: Contradictions at the border. Paper presented at the conference 'Philosophy and Psychology of Vagueness', 21−22 November 2008, École Normale Superieure, Paris. Abstract available at http://paulegre.free.fr/Psyvague/abstract_ ripley.html

Ritov, I. and Baron, J. 1990: Reluctance to vaccinate: Omission bias and ambiguity. *Journal of Behavioral Decision Making*, 3, 263−77.

Scharp, K. 2005: Scorekeeping in a defective language game. *Pragmatics and Cognition*, 13, 203−26.

Serchuk, P., Hargreaves, I. and Zach, R. 2008: Vagueness, logic, and use: Some experimental results. Paper presented at the conference 'Vagueness and Language Use', 7−9 April 2008, École Normale Superieure, Paris. Abstract available at http://paulegre. free.fr/Vagueness/abstracts/serchuk.pdf

Shapiro, S. 2006: *Vagueness in Context*. Oxford: Oxford University Press.

Smith, N. 2009: *Vagueness and Degrees of Truth*. Oxford: Oxford University Press.

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

Sorensen, R. 2001: *Vagueness and Contradiction*. New York: Oxford University Press.

Spranca, M., Minsk, E. and Baron, J. 1991: Omission and commission in judgment and choice. *Journal of Experimental Social Psychology*, 27, 76−105.

Stanley, J. 2003: Context, interest-relativity, and the sorites. *Analysis*, 63, 269−80.

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

Varzi, A. 2007: Supervaluationism and its logics. *Mind*, 116, 663−76.

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

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

Zadeh, L. 1975: Fuzzy logics and approximate reasoning. *Synthese*, 30, 407−28.