

Reconciling Sharp True/False Boundaries With Scalar Vagueness*

Alice I. Kyburg & Lenhart Schubert
University of Rochester

Abstract

The applicability of a vague scalar adjective to an object varies across speakers and contexts. In this paper, we argue that the semantical accounts of vagueness that do not address this variability are inadequate for accounting for truth value intuitions and the information conveyed by utterances. We offer, as a partial replacement, a pragmatic theory of assertability that takes for granted this variation and explains how utterances of sentences with vague scalar adjectives convey information. We propose two bivalent semantic theories of vagueness that bear different relationships to this pragmatic theory. These theories demonstrate that sharp true/false boundaries are reconcilable with scalar vagueness.

Vagueness: An Introduction to the Problem

A word is considered vague in natural language if it allows for borderline cases of applicability. Consider "bald". In a given context, if the number of hairs a man has growing on his head falls within a particular range, then a language user is strongly inclined to: (1) dissent from both the statement *The man is bald* and the statement *The man is not bald*, (2) assent to both statements, or, (3) neither dissent from nor assent to at least one of the statements. Similarly, in a given context, there is a range of heights corresponding to borderline applications of "tall", a range of colors corresponding to borderline applications of "red", and a number range such that a number of grains of sand in a pile that is within that range corresponds to a borderline applications of "heap". Verbs, adverbs, and most other parts of speech can have borderline applications as well.

Vagueness is prevalent throughout natural language dialogue for good reason. Vagueness is of

use in cases where speakers find it most advantageous to their conversational goals *not* to convey the details of their knowledge, e.g., when being precise is too time consuming. It is also useful in cases where speakers are unable to be exact because of an absence of available information. People use vague words regularly in conversation and successfully communicate with sentences containing these words.

Recently, there have been accounts of vagueness proposed in the formal semantics literature that attempt to define truth-conditions for sentences containing vague scalar adjectives that are compatible with speakers' truth-value intuitions. These accounts include truth-value theories and supervaluation theories, both of which are founded on the assumption that if F is a vague predicate, then there is (or could be) a uniquely referring name a such that Fa is neither true nor false. This provision is intended to account for the intuitions people have that some sentences of the form X is A are neither true nor false. In addition, each of these theories involves complex semantic maneuvering to ensure that the truth-values of complex expressions involving borderline cases conform to "philosophical" intuitions concerning logical relations. For example, supervaluation theories require that the semantic value of $Fa \& \neg Fa$ equals 0 (i.e., Falsity) for all vague and non-vague predicates F , and that the semantic value of $Fa \vee \neg Fa$ equals 1 (i.e., Truth) for all vague and non-vague predicates F .

While these complex semantic theories conform to some strong and common intuitions that simpler semantic theories do not conform to, they fail to conform to other equally strong and common intuitions, namely, the "illogical" intuitions of ordinary speakers. In addition, these purely semantic theories are difficult to work into an account of the successful communication that occurs between discourse participants. They do not address the fact that there are at least two individuals involved in ordinary communication, and that these two individuals do not have precisely the same concep-

*This work was supported in part by ONR/DARPA research grant no. N00014-92-J-1512 and by NSF grant no. IRI-9013160.

tion of the applicability criteria for words exhibiting scalar vagueness. We propose a theory of assertability that accounts for some of the intuitions semanticists have attempted to account for and that takes for granted the differing applicability conditions a vague word can have for different dialogue participants and for a given dialogue participant in different instances. With this powerful pragmatic tool, the difficult problems with which the semantic theories of vagueness struggle are more easily addressed. We propose two bivalent semantic theories that work in conjunction with the pragmatic theory to account for these problems. The first semantics allows for considerable discrepancy between truth and assertability. It provides a somewhat unintuitive account of the truth-values of simple and complex sentences while the pragmatic theory provides an independent account of their assertability, from which it is assumed that common intuition is derived. The second semantics assumes a close tie between assertability and truth, and attributes seemingly "illogical" truth-value judgments to delicate context dependencies of the applications of the relevant vague terms. The two semantic theories are offered as demonstrations that sharp true/false boundaries are reconcilable with scalar vagueness.

Truth-Value and Supervaluation Theories

In truth-value approaches to vagueness, the semantic evaluations of sentences are 0, 1, or one or more intermediate values.¹ A truth-value that is not 0 or 1 is called *indefinite*. A common criticism of truth-value approaches is that no intuitive truth-value approach can respect all of the required logical relations and still be truth-functional. For example, to be truth-functional, if *small(x)* is true to the same degree as \neg *red(x)*, then *red(x) & small(x)* must have the same truth-value as *red(x) & \neg red(x)*. It follows that either the truth-value of *red(x) & small(x)* is 0 whenever \neg *red(x)* and *small(x)* are true to the same intermediate degree, or *red(x) & \neg red(x)* has a truth-value other than 0.²

Although this criticism has wide appeal among semanticists, we disagree with one aspect of it. A semantics for natural language that assigns non-zero numeric truth-values to some utterances of sentences of the form *X is A and X is not A* is compatible with widely held truth-value intuitions. It is common practice for ordinary people to accept as true sentences of the form *X is A and X is not A* in the case where *X* is a borderline case of a thing that is *A*. For example, under many

circumstances, a competent speaker would accept as true the sentence *Zeno is tall and Zeno is not tall* if Zeno's height is approximately 5'11". Nevertheless, if the intuitions corresponding to this practice are accepted as semantic data, truth-value accounts have another problem. In the case in which a truth-value account defines the truth-value of a conjunction to be compatible with a speaker's intuitions that a sentence of the form *X is A and X is not A* can be true, this definition is incompatible with the same speaker's intuitions concerning the truth-values of the conjuncts. A truth-value account requires that the truth-value of a conjunction be the minimum of the truth-values of its conjuncts. Yet, a competent speaker who accepts a conjunction of the above form as true holds the intuition that the truth-value of each of the conjuncts is less true than the conjunction, or even false. *Zeno is tall and Zeno is not tall* would each be judged to be less true than *Zeno is tall and Zeno is not tall* in cases where the conjunction were judged to be true.

Supervaluation theories were developed for natural language to address the problems that truth-value theories have with logical relations.³ In this approach, vagueness is understood as lack of specification. The truth-value of a sentence with a vague expression is derived from the truth-values corresponding to each possible way of totally eliminating the vagueness from that sentence. In an early theory proposed by Kamp, the truth-value of a sentence of the form *X is A*, is undefined if some ways of fully precisifying the model include the referent of *X* in the extension of *A* and others include it in the extension of the negation of *A*. While the truth-value of a sentence of the form *X is A* is undefined in cases where the referent of *X* is a borderline case of the property corresponding to *A*, the truth-value of any sentence of the form *X is A and X is not A* is always false since there is never a fully precisified model that includes the referent of *X* in both the positive and negative extension of the predicate corresponding to *A*. Supervaluation theories support other logical relations as well.

Although the truth values assigned by supervaluation theories to complex sentences with logical relations conform to the intuitions of some speakers, they violate the intuitions of other speakers. In particular, since supervaluation theories are designed to assign truth values to complex sentences that support the logical relations, the truth values they assign violate the intuitions described above that conjunctions of contraries with vague scalar adjectives can be true. Thus, with respect to this set of "illogical" intuitions, both truth value theories and supervaluation theories are unsatisfactory.

¹See [Zadeh, 1965] and [Goguen, 1969]

²[Kamp, 1975]

³[Fine, 1975],[Kamp, 1975]

These facts suggest that the complexity of these semantic theories does not lead to sufficient improvement in explanatory adequacy to justify it. We believe that a good approach to take in developing a semantic theory of vagueness is to consider what role the theory might play in a larger theory of communication. From a computational perspective, a theory of meaning that plays no role in a theory of communication is of little use. From a philosophical perspective, a theory of communication may provide clues to a philosophically sophisticated theory of meaning.

One relevant consideration for a theory of communication is whether or not competent speakers in a linguistic community rely on the assumption that other speakers judge words to be applicable (in virtue of their meaning) under exactly the same circumstances as they do. Although speakers may deem non-vague words applicable under identical conditions, it is highly unlikely that individuals share precisely the same conceptions of the applicability criteria for words exhibiting scalar vagueness. Furthermore, speakers do not assume that they do. Consider what would be required for members of a linguistic community to acquire a completely uniform notion of the conditions under which someone would deem "bald" applicable to a man. Our visual judgements of hairiness are far too imprecise and dependent on lighting, visual angle, arrangement of the residual hair, etc., to ensure uniformity from speaker to speaker. Furthermore, if one were to assume that applicability criteria for vague words composed the truth-conditions of sentences with corresponding vague predicates, then a speaker might have to resort to hair-counting to get agreement on the true/false boundary. Leaving a truth value gap, as is posited by 3-valued approaches and supervaluation theories, does not help the situation. Instead of a sharp boundary, between truth and falsity, a gap requires that members of the linguistic community agree on two sharp boundaries: the boundary between truth and indefiniteness, and the boundary between indefiniteness and falsity. Matters are made still worse by having continuously variable truth values. In this case, for the community to share an identical applicability criterion for "bald", it would have to agree on the degree of truth of "bald" for men with n hairs on their heads, for all n . Matters would be aggravated further if factors other than simply the number of hairs, e.g., scalp area, the distribution of hair, etc., needed to be taken into account.

What this suggests is that no amount of semantic maneuvering concerning the meanings of vague terms can eliminate the speaker-to-speaker variation in the applicability criteria for vague words or explain how people succeed in communicating using vague words. A theory of vagueness should take for

granted the variability of applicability criteria and explain how utterances convey information despite it. Furthermore, if semantic maneuvering is not the answer, then it is reasonable to ask whether a simple semantics, such as one that posits the existence of sharp true-false boundaries, is reconcilable with an account of "conveyed information". We believe that it is.

Belief Revision

The pragmatic account of vagueness we propose assumes that the context in which a sentence is uttered contains the information a speaker uses to derive assertability conditions for sentences containing vague words. Although there are many aspects of speaker and hearer cognition that play a role in a speaker's decision to utter a particular sentence, the theory of assertability we propose considers only the speaker's conception of lexical applicability conventions on a given occasion.

Informal Description

Suppose Muriel and Scott are sitting in the faculty lounge discussing the newly appointed faculty member, Zeno. Scott has met the man, but Muriel has not. Muriel has a wide range of expectations about what Zeno might be like. With respect to his height, Muriel positively believes that he cannot have a height outside of the range from 4' to 8'. She has stronger expectations that his height lies in the range between 5'5" and 6'5" than that it lies in any other 12" range. She also has expectations concerning individual heights. She believes that the likelihood that his height is approximately 5'11" is the same as the likelihood that his height is approximately 6'1". She believes that Zeno is most likely to be about 6' tall.

When Muriel asks about Zeno, Scott begins a description of what Zeno looks like with an utterance of the sentence *Zeno is tall*. Muriel assumes that Scott knows approximately what Zeno's height is and that he is speaking sincerely. She also assumes that he is using the word "tall" approximately as she would use the word in this context. She would not deem "tall" applicable to a man whose height was less than 5'10". She would deem "tall" applicable to a man whose height was greater than 6'1". She is unsure of whether she would deem "tall" applicable to a man whose height is between 5'10" and 6'1". Probably, she would not deem "tall" applicable to a man whose height is equal to or a little greater than 5'10", and probably, she would deem "tall" applicable to a man whose height is equal to or a little less than 6'1".

Because Muriel assumes Scott's applicability conventions for "tall" are about the same as her own, the information conveyed to Muriel by his

utterance has predictable consequences on her expectations and beliefs about Zeno's height. Since she assumes in this context that Scott would never deem "tall" applicable to a man whose height is under 5'10", she infers that Zeno's height is definitely greater than or equal to 5'10". Also, since she thinks that there is a greater chance that Scott would deem "tall" applicable to a man who is approximately 6'1" than to a man who is approximately 5'11", and since prior to Scott's utterance her expectations of Zeno being either of these heights were equally strong, she infers from Scott's utterance that the likelihood that Zeno is approximately 6'1" is greater than the likelihood that he is approximately 5'11". Scott's utterance is of absolutely no consequence to her prior belief that Zeno is definitely not taller than 8'.

Suppose that in the place of *Zeno is tall*, Scott had uttered the sentence *Zeno is not tall*. Given the same assumptions about Scott as above, that he knows approximately what Zeno's height is and that he uses words in approximately the same way that Muriel does, the information that would have been conveyed to Muriel by the utterance of the negated sentence would have had clear consequences on her expectations and beliefs. Muriel would deem "not tall" applicable to a man in those cases where she would *definitely not* deem "tall" applicable to him, i.e., in cases where his height is less than 5'10". Similarly, she would *not* deem "not tall" applicable to a man in those cases where she *definitely* would deem "tall" applicable to him, i.e., in cases where his height is greater than 6'1". Probably she would deem "not tall" applicable to a man whose height is equal to or only a little greater than 5'10", and probably, she would *not* deem "not tall" applicable to a man whose height is equal to or only a little less than 6'1". Consequently, following the utterance, Muriel would have revised her original set of beliefs to include the belief that Zeno's height is definitely less than 6'1", and the belief that it is more likely that Zeno's height is approximately 5'11" than that his height is approximately 6'1". She would continue to believe that Zeno's height is definitely not less than 4'.

Finally, suppose that in the place of *Zeno is tall*, Scott had uttered the sentence *Zeno is tall and he's not tall*. Although one might expect Muriel to have found this utterance entirely incoherent, Muriel in fact would have had another common reaction. She would have inferred that Zeno is a borderline case of tall.⁴ Muriel would have expected Scott to utter this conjunction only under the conditions that she herself might have uttered it. Muriel might have

⁴For the purposes of this paper, the reader need not fully embrace this intuition, so long as she acknowledges the fact that there exist a significant number of seemingly competent speakers who hold it.

deemed "tall and not tall" applicable to a man in this context in the cases where the man's height was *not* in a range where she would definitely deem "tall" applicable to him, and it was *not* in a range where she would definitely deem "not tall" applicable to him. That is, she might have deemed "tall and not tall" applicable to a man in this context only if his height were greater than 5'10" and less than 6'1". Consequently, following Scott's utterance of this sentence, Muriel would have revised her original beliefs to include the belief that Zeno's height is definitely in the range between 5'10" and 6'1".

These scenarios are examples of the sort of belief revision that results from normal everyday communication between speakers. Clearly, under normal circumstances, dialogue participants assume that they use their words in about the same way as those with whom they are engaged in dialogue. Hearers make this assumption when they revise their expectations and beliefs following an utterance. One might question whether a hearer really has a huge set of weighted beliefs and expectations concerning every attribute of every object she discusses. An important claim for this theory is that, at the very least, there exists a pool of information available to a hearer from which she can produce these weighted beliefs and expectations. The information conveyed by an utterance modifies this pool of information, and, consequently, changes the hearer's beliefs.

Formal Description

To formally describe the impact an utterance has on a hearer's beliefs, we assume that corresponding to every scalar adjective spoken meaningfully in an utterance is a set of weighted conditional sentences that reflect the information the scalar adjective, as part of that utterance, conveys to the hearer. These sentences correspond to vague and non-vague aspects of a vague scalar adjective. They reflect the hearer's conception of the probability that the speaker would deem a particular word applicable to an object given that the speaker believed the object to have a particular measure of the attribute. In the discussion about Zeno, for instance, some weighted conditional sentences that reflect the *non-vague* information "tall" carries for Muriel in that context include (1) the probability that Scott would deem "tall" applicable to a man who is approximately 5'5" is 0, and (2) the probability that Scott would deem "tall" applicable to a man who is approximately 6'4" is 1. Some weighted conditional sentences that correspond to the *vague* aspects of "tall" that are compatible with Muriel's beliefs are (3) the probability that Scott would deem "tall" applicable to a man who is approximately 5'11" is .5, and (4) the probability that Scott would deem

"tall" applicable to a man who is 6" is .9.

Each weighted conditional sentence corresponding to the information conveyed by a vague scalar adjective in an utterance can be represented as an ordered pair whose first member is the value of the relevant attribute and whose second member is a number which represents the hearer's subjective notion of the likelihood that the speaker would deem the adjective applicable to an object with that attribute value. For example, the weighted conditional sentences listed above can be written:


1. (5'5", 0.0)
2. (6'4", 1.0)
3. (5'11", 0.5)
4. (6'1", 0.9)

The complete set of weighted conditional sentences corresponding to the vague word in an utterance reflects the totality of information conveyed by that word in that utterance. The set is represented as the complete set of corresponding ordered pairs. We will call this set the *compatibility function* for the vague adjective in the utterance.⁵

Consider the compatibility function corresponding to Muriel's conception of Scott's applicability conventions for "tall" in his utterance of *Zeno is tall*. Muriel may have beliefs represented by the following set of weighted conditional sentences:

- ...
- $P(\text{"tall"}|h(x) = 5'8" \pm 1/2") = 0$
 - $P(\text{"tall"}|h(x) = 5'9" \pm 1/2") = 0$
 - $P(\text{"tall"}|h(x) = 5'10" \pm 1/2") = .1$
 - $P(\text{"tall"}|h(x) = 5'11" \pm 1/2") = .5$
 - $P(\text{"tall"}|h(x) = 6'0" \pm 1/2") = .9$
 - $P(\text{"tall"}|h(x) = 6'1" \pm 1/2") = 1$
 - $P(\text{"tall"}|h(x) = 6'2" \pm 1/2") = 1$
- ...


This set of weighted conditional sentences can be represented by a step graph of their corresponding ordered pairs. An idealized representation of Muriel's beliefs about the speaker's word applicability judgements for "tall" is the continuous function C_{tall} :



⁵This term is borrowed from Zadeh, who uses it for a related purpose.

According to the compatibility function, C_{tall} , Muriel believes that Scott would definitely deem "tall" applicable to a faculty member who was 6'1" \pm 1/2" or taller and that Scott would definitely not deem "tall" applicable to a faculty member who is 5'8" \pm 1/2" or shorter. She also believes that likelihood that Scott would deem "tall" applicable to a faculty member who is 6' tall is .9.

In general, a hearer's beliefs concerning the attribute value for a given object is represented as a subjective distribution function. For instance, in the discussion about Zeno, Muriel has probabilistic expectations about the height of the faculty member under consideration. Recall that she believes that Zeno's height is in the range of 4' to 8'. She also believes that the likelihood that his height is approximately 5'11" is the same as the likelihood that his height is approximately 6'1" and that the approximate height Zeno is most likely to measure is 6'. Suppose the following function P best represents her beliefs prior to Scott's utterance:




Equipped with compatibility functions and corresponding subjective distribution functions, the process of revising a hearer's beliefs with information conveyed by an utterance can be formalized. Formally, a distribution function corresponding to a hearer's prior beliefs is combined with a compatibility function corresponding to the information conveyed to the hearer by the utterance to yield a new distribution function which corresponds to the speaker's revised beliefs. One way to combine these functions is by Bayesian updating. The new distribution function is the result of a point by point application of Bayes' rule:

$$P(a|b) = \frac{P(b|a) \cdot P(a)}{P(b)}$$

$P(b|a)$ corresponds to a point on the relevant compatibility function, $P(a)$ corresponds to a point on the prior distribution function, and $P(a|b)$ corresponds to a point on the resulting modified distribution function. The probability that an object has a particular measure of a scalar property given that the speaker would deem ψ applicable to that object is the product of the probability that the speaker would deem ψ applicable to that object given that it has that particular measure and the speaker's prior degree of expectation that the object has that par-

ticular measure, divided by the likelihood that the speaker would deem ψ applicable (whatever measure it has).

After Scott utters *Zeno is tall*, Muriel's new expectations and beliefs concerning Zeno's height are reflected in an updated distribution function. This function is derived from (1) the prior subjective distribution P reflecting her prior expectations of Zeno's height, and (2) the compatibility function C_{tall} corresponding to Muriel's conception of the Scott's applicability conventions for "tall" in this context. For example, if Muriel's conception of the applicability of the adjective "tall" to a man whose height is approximately 6'0" is .9, Muriel can use this information in combination with her subjective notions of how likely it is that Zeno would have a height of approximately 6'0" (prior to being described with "tall") to arrive at the new likelihood that he is approximately 6'0". If this procedure is applied to every height, then the result will be the updated distribution function Q , corresponding to Muriel's revised beliefs.




A formal account of the hearer's change in expectations and beliefs upon hearing an utterance of the form X is A and B (taken as an abbreviation for X is A and X is B) results if the conjunction is assumed to be processed as if the speaker made two consecutive utterances, corresponding to X is A and X is B , and they were simply processed in succession according to the belief revision process just described. The distribution function representing the hearer's prior expectations and beliefs is updated with the compatibility function corresponding to the information conveyed by A in this utterance, yielding an intermediate distribution function. This intermediate function, in turn, is updated with the compatibility function corresponding to the information conveyed by B in this utterance, yielding the distribution function which represents the hearer's expectations and beliefs following the speaker's utterance of the conjunction.

When the vague scalar adjectives in a conjunction correspond to the same scalar property, this algorithm is mathematically equivalent to updating the prior distribution function with the product of the two compatibility functions. $P((h(x) = i) | (A \text{ and } B \text{ are deemed applicable to } x)) =$

$$\frac{C_A(h(x) = i) \cdot C_B(h(x) = i) \cdot P(h(x) = i)}{P(A \text{ and } B \text{ are deemed applicable to } x)}$$

A common case of different scalar adjectives corresponding to exactly the same attribute is the case where one is the negation of the other. Technically, the latter is not an adjective, of course, but the compatibility function corresponding to the negation of a vague adjective conveys analogous information as that corresponding to the vague adjective.

Consider Scott's utterance of *Zeno is tall and he's not tall*. Given the prior distribution function P corresponding to Muriel's prior expectations of Zeno's height, and given the compatibility functions C_{tall} and $C_{not-tall}$ for the utterance, the result of revising Muriel's prior expectations to take into account the information conveyed by this utterance would correspond to P , first updated by the information conveyed by the utterance of *Zeno is tall* to yield Q , and then updated with the information conveyed by the utterance of *Zeno is not tall*, to yield R , where R is the distribution function that corresponds to Muriel's revised beliefs following Scott's utterance of *Zeno is tall and he's not tall*.



The information conveyed to Muriel by the utterance of the conjunction is reflected in the compatibility function $C_{tall \text{ and } not-tall}$, which is the point by point product of C_{tall} and $C_{not-tall}$. Both intuition and preliminary experimental results suggest that the results this formula yields correspond to the revised beliefs hearers actually have upon processing utterances of this kind.⁶ But Bayes' rule, applied to a formula of the form $P(a|b\&c)$, yields $P(a) \cdot P(b|a) \cdot P(c|a) / P(b\&c)$ if and only if b and c are independent variables. Clearly, the hearer's conception of the applicability of the negation of A to X in the utterance context is dependent on her conception of the applicability of A to X in that context. Yet, Muriel's new beliefs about Zeno's height can be derived from her old beliefs and the probabilities corresponding to the two dependent sets of beliefs that reflect Muriel's conception of the applicability of "tall" and the applicability of "not

⁶[Schubert]