

# Vagueness, Probability, and Linguistic Representation

## 1. Introduction

Intuitively, vague words are words which clearly apply to certain objects, clearly do not apply to other objects, and which have certain intermediate cases, where it is unclear whether the term applies or not. Standard examples are scalar adjectives like “bald”, “tall”, and “fat” and common nouns such as “heap”. Vagueness is a serious problem for the logical analysis of language: the classical statement of this problem is the *sorites paradox*. One grain of sand is clearly not a heap. It seems plausible that, if you have something that is not a heap and you add one grain of sand to it, you still do not have a heap. But from these two premises it follows that no amount of sand can constitute a heap. That is, the following is a valid argument:

### (1) Sorites Paradox

- a. One grain of sand is not a heap.
- b. If you add one grain of sand to something that is not a heap, then you still will not have a heap.
- 
- c. No pile of sand, no matter how large, is a heap.

The problem of the sorites, in a nutshell, is that the second premise (the “inductive premise”) is intuitively plausible, but the conclusion is not. We know that the first premise is true (one grain is not a heap), and that the conclusion is false (heaps of sand do exist). It follows that the inductive premise is false. But it is difficult to find an intuitively plausible and logically satisfying theory of where the inductive premise goes wrong.

One prominent approach to the problem of the sorites is the epistemic theory (Williamson 1994). According to this theory, meanings *are* precise, and the phenomena of vagueness are the result of humans’ imperfect knowledge of the meanings of words. If this is right, vagueness is not a *logical* problem, and semantics can carry along merrily, leaving the explanation of the sorites to epistemologists and psychologists. However, a common objection to the epistemic theory is that it requires an implausible divorce between meaning and humans’ knowledge and use of language (e.g., Lassiter 2008). I find this aspect of epistemicism undesirable. However, I also think that the epistemic theory contains a crucial insight which I try to recapture in this work: vagueness is not located in the semantic theory *per se*, but in the relation between human agents and the semantic theory.

The theory of vagueness described in this paper incorporates this aspect of the epistemic theory while also relating meaning to the linguistic knowledge and behavior of speakers and listeners. It can be seen as a development of Lewis’ (1972) suggestion that “languages themselves are free of vagueness but ... the linguistic conventions of a population, or the linguistic habits of a person, select not a point but a fuzzy region in the space of precise languages”. Lewis, like epistemic theorists, envisions a theory in which languages are precise and the underlying logic is bivalent. The difference is that Lewis does not endorse the epistemicist’s claim that there is a single language that is being spoken in a given conversation. Rather, there is always a range of languages that are

contenders for being the language of conversation, and the epistemic problem is to use prior knowledge and context to select the most plausible candidates for being the language of conversation.

This paper attempts to capture the spirit of Lewis' proposal within a probabilistic theory of semantic representation. The main aspect of the probabilistic theory of vagueness which I will explore here is its treatment of the sorites paradox. However, one of the most interesting aspects of the probabilistic approach, which will remain implicit in the following discussion for reasons of space, lies in the connection it draws between vagueness and other aspects of human cognition. There is considerable evidence from psychology that human learning involves induction over weighted alternatives (Tenenbaum 1999; Xu and Tenenbaum 2007) and that belief and reasoning also crucially incorporate probabilistic information (e.g., Chater and Oaksford 2008). Sociolinguists have long made use of probabilistic representations (Weinreich, Labov and Herzing 1968). More recent trends in theoretical linguistics suggest that grammatical knowledge is also uncertain, and that probabilistic models are useful in linguistics as well (Yang 2002; Bod et al. 2003). The main innovation of the present work is to extend this approach from grammar to semantic representation, and to show that a plausible solution to the sorites paradox follows. Thus, if the theory is correct, then vagueness is not a specifically semantic phenomenon, but a consequence of the nature of linguistic knowledge and general principles of language learning and use.

The paper is structured as follows: section 2 describes Robert Stalnaker's theory of assertion and his argument for a unified possible-worlds/languages model of belief and assertion. A probabilistic modification of this model is described and illustrated in section 3. In section 4 the model is applied to the interpretation and representation of vague terms. Finally, in section 5 I show that this model gives us an explanation of the sorites paradox which explains not only why the paradox is invalid, but also why it seems so compelling.

## **2. Metalinguistic assertion and linguistic belief**

In Stalnaker's (1978) theory of assertion, the role of an assertion is to eliminate certain possibilities from the common ground, which is construed as a set of worlds considered by the conversational participants as live possibilities for how the actual world might be. Suppose you don't know whether it is raining outside. If a reliable source tells you, "It is raining", you will normally change your beliefs so that you no longer consider worlds in which it is not raining to be candidates for being the actual world.

Stalnaker's theory of assertion is not restricted to modeling update of non-linguistic beliefs, however. Among the assumptions that speakers bring to a conversation are beliefs about the linguistic context, such as a prior theory about what kinds of sounds are likely to be useful in communicating information to a given audience. As Stalnaker (1978, 2004) and Barker (2002) show, in addition to eliminating possible ways the non-linguistic world might be from the common ground, assertions may also eliminate possible languages from the common ground.

Suppose someone asks you, 'What is an optometrist?' Imagine that your interlocutor's state of knowledge is such that there are two sets of possibilities. Let  $L_1$  designate all languages in which *optometrist* and *eye doctor* are mapped to the same concept or set of individuals (whichever your preferred semantic theory treats).  $L_2$

designates all languages in which *optometrist* and *plumber* are mapped to the same concept or set of individuals. In this context, the reply ‘An optometrist is an eye doctor’ serves to eliminate  $L_2$  from the common ground. The net effect is to inform the linguistically uncertain interlocutor that, in the current language of conversation, the sequence of sounds *optometrist* is not an appropriate way to communicate any concept other than EYE DOCTOR. Importantly, *An optometrist is an eye doctor* gives no information about the non-linguistic world, but rather functions as an instruction to interpret a certain noise in a certain way.

The upshot of Stalnaker’s approach to metalinguistic discourse, then, is a unified model of assertion using two parallel but formally similar dimensions of possible languages and possible worlds.

### 3. Probabilistic beliefs, assertion and inference

In section 2 we saw evidence from the phenomenon of metalinguistic assertion that linguistic behavior depends on and makes reference to speakers’ and listeners’ linguistic beliefs, and potentially affects their subsequent linguistic behavior. On the other hand, there is considerable psychological and linguistic evidence that beliefs in general and linguistic representations in particular are probabilistic in nature, as mentioned above. We can integrate these insights by representing an agent’s beliefs as a set of ordered pairs each consisting of a world and a real number in the range  $[0,1]$ , subject to the condition that the sum of the probabilities attached to all worlds in the agent’s belief-set is 1. This model is readily implemented in Kripke semantics and is logically well-behaved: see Halpern 1997 for an axiomatization with soundness and completeness proofs.

Suppose that an agent A considers four worlds possible in some context C.

- (2) In  $w_1$  it is snowing and John is a plumber.
- In  $w_2$  it is not snowing and John is a plumber.
- In  $w_3$  it is snowing and John is an eye doctor.
- In  $w_4$  it is not snowing and John is an eye doctor.

A’s belief-set might look like (3):

- (3)  $B_A = \{ \langle w_1, .1 \rangle, \langle w_2, .4 \rangle, \langle w_3, .3 \rangle, \langle w_4, .2 \rangle \}$

We can calculate A’s subjective probability of any proposition  $p$  by summing the probabilities of all worlds in which  $p$  is true. So, for example,  $p_A(\text{it is snowing}) = p_A(w_1) + p_A(w_3) = 0.4$ . Similarly,  $p_A(\text{John is a plumber}) = p_A(w_1) + p_A(w_2) = 0.5$ .

Possible languages can be treated in similar fashion as probability distributions over possible languages. On this approach, a context set consists of non-linguistic beliefs – modeled as a set of possible world-probability pairs – and metalinguistic beliefs – modeled as a set of possible language-probability pairs.

- (4) *Possible languages* are functions from utterances to model-theoretic objects.

For example, suppose that an agent A considers four languages to be viable candidates for being the actual language of conversation in the context given above.

- (5) In  $L_1$ , ‘optometrist’ and ‘plumber’ mean the same, and ‘I’ picks out the **speaker**.  
 In  $L_2$ , ‘optometrist’ and ‘eye doctor’ mean the same, and ‘I’ picks out the **speaker**.  
 In  $L_3$ , ‘optometrist’ and ‘plumber’ mean the same, and ‘I’ picks out the **listener**.  
 In  $L_4$ , ‘optometrist’ and ‘eye doctor’ mean the same, and ‘I’ picks out the **listener**.

The revised representation of A’s belief-state is a tuple of non-linguistic and linguistic beliefs, as in (3’):

$$(3') \quad B_A = \langle \{ \langle w_1, .1 \rangle, \langle w_2, .4 \rangle, \langle w_3, .3 \rangle, \langle w_4, .2 \rangle \}, \\ \{ \langle L_1, .3 \rangle, \langle L_2, .4 \rangle, \langle L_3, .1 \rangle, \langle L_4, .2 \rangle \} \rangle$$

Again we calculate the subjective probability of an interpretation by summing the probabilities of all languages which yield that interpretation. So,  $p_A(\text{‘optometrist’} = \text{‘plumber’}) = p_A(L_1) + p_A(L_3) = .4$ .

Now suppose that A comes to learn that John is an eye doctor. This will affect non-linguistic beliefs, but not linguistic beliefs; and so (5) will be updated so that the probability of all worlds in which John is not an eye doctor is 0.<sup>1</sup>

$$(6) \quad B_A = \langle \{ \langle w_1, 0 \rangle, \langle w_2, 0 \rangle, \langle w_3, .6 \rangle, \langle w_4, .4 \rangle \}, \\ \{ \langle L_1, .3 \rangle, \langle L_2, .4 \rangle, \langle L_3, .1 \rangle, \langle L_4, .2 \rangle \} \rangle$$

Now suppose John says to A:

- (7) I am an optometrist.

In  $L_1$ , (7) means that John is a plumber. In  $L_2$ , (7) means that John is an eye doctor. In  $L_3$ , (7) means that A is a plumber; in  $L_4$ , (7) means that A is an eye doctor. If A is neither an eye doctor nor a plumber, and he assumes that John is well-informed and cooperative, he will eliminate  $L_3$  and  $L_4$  from consideration. This is because  $L_3$  and  $L_4$  entail that John is uttering a known falsehood. But in addition, A can eliminate  $L_1$  from his belief-set, because he knows that John is an eye-doctor and not a plumber, and accepting  $L_1$  would mean that John is uttering a known falsehood. As a result, (7) – which is on face a statement about the world – does not add non-linguistic information, but it allows A to infer information about the language of conversation. So, after (7) is uttered, A’s belief-set will be as in (8):

$$(8) \quad B_A = \langle \{ \langle w_3, .4 \rangle, \langle w_4, .6 \rangle \}, \{ \langle L_2, 1 \rangle \} \rangle$$

---

<sup>1</sup> The probabilities I give for  $w_3$  and  $w_4$  in (6) should technically be derived using Bayes’ theorem, but it is not necessary to go into this issue for this example: since half the worlds have been eliminated and the fact that John is an eye doctor is irrelevant to the choice between  $w_3$  and  $w_4$ , we simply double the prior probabilities of these worlds.

This simple example illustrates the interaction of non-linguistic and linguistic information in interpretation and belief update.

#### 4. Representing vague terms

Toy examples like the one given above obscure the fact that the domain of *a priori* possible states of the world is extremely large, perhaps infinite. If any function that satisfies (4) counts as a possible language, then possible languages are similar to possible worlds in this respect as well: there is a very large number of possible languages, many of which are minimally different from each other. Normally most of these will be out of contentions, but even native speakers consider various interpretive options as viable in ordinary situations.

Suppose that a speaker A utters some noise U, and an interpreter B takes U to pick out some portion of an ontologically continuous object or property (such as height or the color spectrum). U could in principle designate an unlimited number of distinct properties  $P_1 \dots P_n$  each of which divides the continuum in a different way. Each of these properties, in turn, can be the value of U in a range of distinct possible languages, which can be as similar as you like except for the value of U. So, for example, there are possible languages in which the dividing line between ‘tall’ and ‘not tall’ is 6’0”, 6’1”, 6’1.1”, 6’1.11”, and so on.

The intuitive idea behind the present approach is that vagueness consists in this feature of linguistic knowledge: in any given interaction there is a range of possible languages that we might be speaking, and thus a range of possible meanings for a vague expression like ‘tall’ that might be relevant in a context. But not all interpretations of an utterance are equally plausible: when a predicate is vague there is no point at which a possible language  $L_1$  yields a plausible interpretation of “tall”, and a neighboring possible language  $L_2$  which resolves “tall” in a similar fashion yields an implausible interpretation. Rather, the plausibility of neighboring interpretive theories changes gradually.

This idea can be explicated in the Stalnakerian possible-languages model as follows. As above, the second coordinate of an agent’s belief-set is a set of pairs of a possible language and a real number in the range  $[0,1]$ . We extract from the agent’s belief-set the representation of an utterance in the following way. Let **PL** be the set of possible languages, whose members are  $L_1, L_2, \dots L_n$ . As usual, **D<sub>e</sub>** is the set of possible objects whose members are **o<sub>1</sub>, o<sub>2</sub>, ... o<sub>n</sub>**. For simplicity’s sake we restrict attention to model-theoretic objects of type  $\langle e, t \rangle$  (e.g. common nouns and scalar adjectives), although the definition could easily be modified to allow objects of arbitrary semantic type.

- (9) The *lexical representation*  $LR_{u,A}$  of an utterance  $u$  according to an agent  $A$  is a set of pairs  $\langle \mathbf{o}_m, d \rangle$  of a possible object  $\mathbf{o}_m \in \mathbf{D}_e$  and a real number  $d$  in the range  $[0,1]$ .

a.  $\langle \mathbf{o}_m, d \rangle \in LR_{u,A}$  iff  $d = \sum_{L \in \mathbf{PL}} [p_A(L): (L(u))(\mathbf{o}_m)=1]$

b. If  $\langle \mathbf{o}_m, d \rangle \in LR_{u,A}$  then the subjective probability according to  $A$  that  $u$  applies to  $\mathbf{o}_m$  is  $d$ .

c. I will henceforth abbreviate clause (9b) as:  $p_A(u(\mathbf{o}_m)) = d$ .

((9c) is an abbreviation for readability's sake: taken literally, it would require that an utterance take a model-theoretic object as an argument, which is not possible in since utterances are not functions.)

Clause (9a) stipulates that  $A$ 's subjective probability that an utterance  $u$  applies to an object  $\mathbf{o}_m$  is equal to the sum of the probabilities of all possible languages in  $A$ 's belief-set in which the value of  $u$  applied to  $\mathbf{o}_m$  returns 1. In this way,  $A$ 's probabilistic belief-set determines the probabilistic lexical representation of any possible utterance  $u$ . We can use (9) to translate between possible-languages talk and talk of lexical representations without loss of information. (A similar definition would do the same for possible worlds and propositions in our probabilistic model.)

As an example, suppose that the lexical representation of *tall* for a particular speaker  $A$ ,  $LR_{tall,A}$ , has the following form.  $A$  considers possible these resolutions of *tall*: 5'5", 5'6", ... 6'5" (spaced at 1" to simplify the model). Letting italics represent utterances as above and boldface indicate model-theoretic objects, we'll call these **tall<sub>1</sub>**, **tall<sub>2</sub>**, ... **tall<sub>13</sub>**. Each resolution of *tall* denotes (the characteristic function of) a set of individuals who the conditions abbreviated as "Threshold". For example,  $[[\mathbf{tall}_1]] = \{x: x\text{'s height} \geq 5'5''\}$ ;  $[[\mathbf{tall}_2]] = \{x: x\text{'s height} \geq 5'6''\}$ ; and so forth. The probability distribution in the bottom row of (10) assigns a probability in the range  $[0,1]$  to each resolution of *tall*. For example, the fifth column of (10) is read: "The probability that the denotation of *tall* is  $\{x: x\text{'s height is at least } 5'8''\}$  is 0.03".

(10) Name	<b>tall<sub>1</sub></b>	<b>tall<sub>2</sub></b>	<b>tall<sub>3</sub></b>	<b>tall<sub>4</sub></b>	<b>tall<sub>5</sub></b>	<b>tall<sub>6</sub></b>	<b>tall<sub>7</sub></b>	<b>tall<sub>8</sub></b>	<b>tall<sub>9</sub></b>	<b>tall<sub>10</sub></b>	<b>tall<sub>11</sub></b>	<b>tall<sub>12</sub></b>	<b>tall<sub>13</sub></b>
Threshold	5'5	5'6	5'7	5'8	5'9	5'10	5'11	6'0	6'1	6'2	6'3	6'4	6'5
$p_A([[tall]] = \mathbf{tall}_n)$	0	0	.01	.03	.09	.14	.23	.23	.14	.09	.03	.01	0

As in (9), for any utterance  $U$  which denotes a set of individuals, the probability that  $x$  is  $U$  is the sum of the probabilities of all resolutions of which  $x$  is an element.

Because all available resolutions of *tall* are upward monotonic it is simple to calculate to calculate the probability that an individual of arbitrary height will count as *tall*. Let's consider a series of 13 individuals in this height range, spaced at 1", called  $x_1$ ,  $x_2$ , ...  $x_{13}$ . Using (10) we calculate for each individual  $x_n$  the probability that  $x_n$  is *tall* by summing the probability that  $[[tall]] = \mathbf{tall}_n$  for all  $m \leq n$ . For example,  $p(tall(x_5)) = \sum_{i=1 \dots 5} p([[tall]] = \mathbf{tall}_i) = 0+0+.01+.03+.09 = .13$ . The table in (11) gives the values of  $tall(x_n)$  for a representative sample of individuals of various heights, using (9) and (10).



of **tall** as a model-theoretic object. **Tall** in itself does not represent any object in my system: a subscript is needed to indicate which possible language (function from utterances to model-theoretic objects) **tall<sub>n</sub>** was generated by. So, if *tall* is assigned a value in  $L_{34}$ , then  $L_{34}(tall)$  will be written **tall<sub>34</sub>**. On this interpretation the paradox appears as in (14):

$$\begin{aligned}
 (14) \quad & \text{a. } \neg(\mathbf{tall}_{34}(x_1)) \\
 & \text{b. } \forall n [ (\neg(\mathbf{tall}_{34}(x_n)) \rightarrow (\neg(\mathbf{tall}_{34}(x_{n+1})))) ] \quad \text{(equivalently,} \\
 & \quad \text{-----} \quad \neg \exists n (\neg(\mathbf{tall}_{34}(x_n)) \& (\mathbf{tall}_{34}(x_{n+1}))) ] \\
 & \text{c. } \forall n [\neg(\mathbf{tall}_{34}(x_n))] \\
 & \therefore \neg \mathbf{tall}_{34}(x_{13})
 \end{aligned}$$

Since possible languages are perfectly precise, the inductive premise is plainly false: the bivalence of the semantic metalanguage guarantees that there is a precise cut-off between **tall<sub>34</sub>** and  $\neg \mathbf{tall}_{34}$ , and so the conclusion does not follow and the paradox does not arise.

Suppose now we rewrite the paradox using utterances in place of model-theoretic objects. *Tall* is an utterance, and it is interpreted by some possible language  $L_n$  as a semantic object **tall<sub>n</sub>**. Crucially, *tall* is not in itself a model-theoretic object. If we attempt to restate the paradox using *tall* in the place of **tall**, we get (15).

$$\begin{aligned}
 (15) \quad & \text{a. } \neg(tall(x_1)) \\
 & \text{b. } \forall n [ (\neg(tall(x_n)) \rightarrow (\neg(tall(x_{n+1})))) ] \\
 & \quad \text{-----} \\
 & \text{c. } \forall n [\neg(tall(x_n))] \\
 & \therefore \neg tall(x_{13})
 \end{aligned}$$

None of the clauses in (15) are syntactically or semantically well-formed within the theory I have introduced. I have occasionally spoken of “the probability that *u* applies to *o* (according to *A*)”, but this was explicitly introduced in (9a) as an abbreviatory convention: If  $\langle \mathbf{o}_m, d \rangle \in LR_{u,A}$ , we may say that  $p_A(u(\mathbf{o}_m)) = d$ . However, the bare claim that  $x_n$  is *tall* is meaningless within this theory: we can only say that the utterance *tall* applies to  $x_n$  with some probability  $d$ .

Suppose we rewrite the paradox using probabilities, as the present approach demands. It seems plausible that “ $x_n$  is not tall” should be expressed as “ $p_A(tall(x_1)) = 0$ ”. If we accept this, the restatement of the sorites paradox is:

$$\begin{aligned}
 (15) \quad & \text{a. } p_A(tall(x_1)) = 0 \\
 & \text{b. } \forall n [p_A(tall(x_n)) = 0 \rightarrow p_A(tall(x_{n+1})) = 0] \\
 & \quad \text{-----} \\
 & \text{c. } \forall n [p_A(tall(x_n)) = 0] \\
 & \therefore p_A(tall(x_{13})) = 0
 \end{aligned}$$

(15) is logically valid, but premise (b) is much less intuitively plausible than the original inductive premise (13b). There is simply no reason to assume that, if the probability that



something is *tall* is 0, the probability that an adjacent item is *tall* must also be 0 (rather than some small but non-zero amount).

Much more plausible is the probabilistic translation of the existential variant of the inductive premise, “ $\neg \exists n [(\neg(\mathbf{tall}(x_n)) \rightarrow (\neg(\mathbf{tall}(x_{n+1}))))]$ ”. A reasonable translation is (16):

$$(16) \quad \neg \exists n [p_A(\mathbf{tall}(x_n)) = 0 \ \& \ p_A(\mathbf{tall}(x_{n+1})) = 1]$$

But (16) is not equivalent to (15b) in the current system: denying that there is a point at which the probability function jumps from 0 to 1 is not the same as denying that it ever increases from 0. (16) is true for *tall* and any other vague predicate, but this creates no problem: (16) is, if anything, just a necessary condition for a predicate’s being vague.

Fara (2000) suggests that a convincing theory of the sorites, if it denies the inductive premise, must answer three separate questions (slightly modified from Fara 2000).

**1. The Semantic Question:** If (a) is not true, then must this classical equivalent of its negation, the “sharp boundaries” claim, be true?

The “sharp boundaries” claim:  $\exists n [\neg \mathbf{tall}(x_n) \ \& \ \mathbf{tall}(x_{n+1})]$

(a) If the sharp boundaries claim *is* true, how is its truth compatible with the fact that vague predicates have borderline cases? For the sharp boundaries claim seems to deny just that.

(b) If the sharp boundaries claim is *not* true, then given that a classical equivalent of its negation is not true either, what revision of classical logic and semantics must be made to accommodate that fact?

**2. The Epistemological Question:** If “ $\exists n [\neg \mathbf{tall}(x_n) \ \& \ \mathbf{tall}(x_{n+1})]$ ” is not true, why are we unable to say which one (or more) of its instances is not true—even when all the heights of the possible values of “ $x_n$ ” are known?

**3. The Psychological Question:** If inductive premise is not true, why were we so inclined to accept it in the first place? In other words, what is it about vague predicates that makes them seem tolerant, and hence boundaryless to us?

Let’s address these questions in turn.

### 1. The semantic question.

(a) If we replace **tall** in Fara’s formulation of the “sharp boundaries” claim by a model-theoretic object acceptable in our system such as **tall<sub>n</sub>**, the claim is true. This is not problematic because our original intuition that the sharp boundaries claim is false, and that the universal sorites premise is true, was not an intuition about some model-theoretic object **tall<sub>n</sub>** but an intuition about the meaning of the word (utterance) *tall*.

(b) If we replace **tall** in Fara’s formulation of the “sharp boundaries” claim by an utterance such as *tall*, making appropriate adjustments (as in (16)), the “sharp boundaries” claim is false. However, no revision of classical logic and semantics is required to explain these facts; rather, this result follows from the fact that the utterance *tall* does not denote a unique object, but denotes various objects with differing probabilities. The semantic metalanguage is nevertheless classical.

**2. The epistemological question.** “ $\exists n [\neg \mathbf{tall}(x_n) \ \& \ \mathbf{tall}(x_{n+1})]$ ” is not well-formed

in the present theory. If we substitute **tall<sub>n</sub>**, as in “ $\exists n [\neg \text{tall}_n(x_n) \ \& \ \text{tall}_n(x_{n+1})]$ ” we can identify which *n* satisfies this formula given a complete specification of the language *L<sub>n</sub>* or of the extension of **tall<sub>n</sub>**. If we consider the sharp boundaries claim substituting the utterance *tall*, our language does not permit us to ask which *n* satisfies “ $\exists n [\neg \text{tall}(x_n) \ \& \ \text{tall}(x_{n+1})]$ ”, because this sentence is not well-formed. The approximation of this formula in the present language would be the negation of (16),  $\exists n[p_A(\text{tall}(x_n)) = 0 \ \& \ p_A(\text{tall}(x_{n+1})) = 1]$ . This is, of course, false of any vague predicate.

**3. The psychological question.** I suggest that we are inclined to accept the inductive premise because we interpret it as a claim about words/utterances rather than about model-theoretic objects (which are probably not accessible to introspection anyway, like most grammatical objects). Speakers know that, given a pair of very similar objects, vague words like *tall* will not apply to one with probability 0 and to the other with probability 1. There is of course much more to be said here: in particular, the relationship between probability of application and assertibility needs clarification. Nevertheless, the present theory yields a plausible approach to Fara’s psychological question which seems to me to deserve serious consideration.

## 6. Conclusion

The theory of vagueness described here has an important advantage over many competitors: it stipulates no special semantic apparatus for vague terms, nor does it rely on the claim that the meanings of words are defined independent of speakers’ knowledge of language. Rather, the present theory relies on general and independently motivated properties of language use and human cognition, motivated by philosophical concerns, linguistic data, and evidence that probability plays a role in the representation of linguistic and non-linguistic knowledge. The model is readily implemented in probabilistic modal logic, and yields an account of the lexical representation of vague words. The theory forces us to reinterpret the traditional sorites paradox in one two ways. On one of these the denial of the inductive premise is harmless, and on the other the inductive premise is no longer plausible. Finally, these results are explanatory with respect to Fara’s (2000) three questions for an account of the sorites which denies the inductive premise.

## References

- Barker, Chris. 2002. The dynamics of vagueness. *Linguistics and Philosophy* 25: 1-36.
- Bod, Rens, Jennifer Hay and Stephanie Jannedy (eds.). 2003. *Probabilistic Linguistics*. MIT Press.
- Chater, Nick, and Mike Oaksford (eds.). 2008. *The Probabilistic Mind: Prospects for Bayesian Cognitive Science*. Oxford: Oxford University Press.
- Fagin, Ronald, and Joseph Y. Halpern. 1994. Reasoning about knowledge and probability. *Journal of the Association for Computing Machinery* 41(2): 340-367.
- Fara, Delia Graff. 2000. Shifting sands: an interest-relative theory of vagueness. *Philosophical Topics* 28(1).
- Fine, Kit. 1975. Vagueness, truth and logic. *Synthese* 30: 256-300.
- Halpern, Joseph Y. 1997. A logical approach to reasoning about uncertainty. In X.

- Arrozala, K. Korta, and F. J. Pelletier eds., *Discourse, Interaction, and Communication*. Dordrecht: Kluwer.
- Kamp, Hans. 1975. Two theories about adjectives. In E. Keenan (ed.), *Formal Semantics of Natural Language*. Cambridge: Cambridge University Press.
- Lassiter, Daniel. 2008. Semantic externalism, language variation, and sociolinguistic accommodation. *Mind and Language* 23(5).
- Lewis, David. 1972. General semantics. In D. Davidson and G. Harman eds., *Semantics of Natural Language*. Dordrecht: Reidel. 169-218.
- Lewis, David. 1973. *Counterfactuals*. Cambridge, MA: Harvard University Press.
- Stalnaker, Robert. 1978. Assertion. In Peter Cole ed., *Syntax and Semantics* 9. New York: Academic Press.
- Stalnaker, Robert. 2004. Assertion revisited: On the interpretation of two-dimensional modal semantics. *Philosophical Studies* 118: 299-322.
- Tenenbaum, Joshua. 1999. Bayesian modeling of human concept learning. *Advances in Neural Information Processing Systems 11*, ed. M.S.Kearns, S.A.Solla, & D.A.Cohn. MIT Press.
- Weinreich, Uriel, Labov, William, & Herzog, Marvin. 1968. Empirical foundations for a theory of language change. In W. P. Lehmann & Y. Malkeil (ed.), *Directions for historical linguistics: A symposium*. Austin: University of Texas Press.
- Williamson, Timothy. 1994. *Vagueness*. London: Routledge.
- Williamson, Timothy. 1999. "Schiffer on the Epistemic Theory of Vagueness." *Philosophical Perspectives 13: Epistemology*, ed. James E. Tomberlin. Cambridge, MA: Blackwell. Pp. 505-17.
- Xu, F. and Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological Review* 114(2).
- Yang, Charles. 2002. *Knowledge and Learning in Natural Language*. Oxford University Press.
- Zadeh, Lofti. 1975. Fuzzy Logic and Approximate Reasoning. *Synthese* 30:407-28.