



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Vagueness as probabilistic linguistic knowledge

Citation for published version:

Lassiter, D 2011, Vagueness as probabilistic linguistic knowledge. in R Nouwen, R Rooij, U Sauerland & H-C Schmitz (eds), *Vagueness in Communication: International Workshop, VIC 2009, held as part of ESSLLI 2009, Bordeaux, France, July 20-24, 2009. Revised Selected Papers.*. 1 edn, Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), vol. LNAI, 6517, Springer-Verlag Berlin Heidelberg, Berlin, Heidelberg, pp. 127-150, 21st European Summer School in Logic, Language and Information, Bordeaux, France, 20/07/09.
https://doi.org/10.1007/978-3-642-18446-8_8

Digital Object Identifier (DOI):

[10.1007/978-3-642-18446-8_8](https://doi.org/10.1007/978-3-642-18446-8_8)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Vagueness in Communication

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Vagueness as Probabilistic Linguistic Knowledge

Daniel Lassiter

New York University

To appear in R. Nouwen, U. Sauerland, H.-C. Schmitz and R. van Rooij eds.,
Vagueness in Communication, Springer.

Abstract. Consideration of the *metalinguistic* effects of utterances involving vague terms has led Barker [1] to treat vagueness using a modified Stalnakerian model of assertion. I present a sorites-like puzzle for factual beliefs in the standard Stalnakerian model [28] and show that it can be resolved by enriching the model to make use of probabilistic belief spaces. An analogous problem arises for metalinguistic information in Barker’s model, and I suggest that a similar enrichment is needed here as well. The result is a probabilistic theory of linguistic representation that retains a classical metalanguage but avoids the undesirable divorce between meaning and use inherent in the epistemic theory [34]. I also show that the probabilistic approach provides a plausible account of the sorites paradox and higher-order vagueness and that it fares well empirically and conceptually in comparison to leading competitors.

Keywords: Vagueness, probability, lexical representation, higher-order vagueness

1 Introduction

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. $\therefore$ No pile of sand, no matter how large, is a heap.

Slightly more formally, taking S_n to mean “an aggregation of n grains of sand”:

- (2) a. $\neg(heap(S_1))$
 - b. $\forall n[\neg heap(S_n) \rightarrow \neg heap(S_{n+1})]$
 - c. $\therefore \forall n[\neg heap(S_n)]$

In classical logic, the inductive premise (2b) is equivalent to a denial that there is any number n such that n grains of sand do not constitute a heap, but $n + 1$ grains do. That is,

- (3) $\neg \exists n[\neg heap(S_n) \wedge heap(S_{n+1})]$

The problem of the sorites is that both variants of the inductive premise, (2b) and (3), are 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 very difficult to determine precisely where the inductive premise goes wrong.

In addition to the problem of the sorites, the existence of vagueness presents a serious challenge to the foundations of formal semantics. One way to identify this problem is in Grim’s [8] claim that a precise theory of vague terms is impossible:

Any successful account of vagueness will have to incorporate vagueness in one way or another; at the core of the Supervaluational approach, for example, lies the vagueness of ‘acceptable precisifications’. Any hope for a fully precise account of vagueness is doomed.

A basic assumption of formal semantics is that natural language meaning and inference can be modeled in a mathematically precise fashion. This assumption is embodied directly in the principle of COMPOSITIONALITY: the meaning of an expression is built up from the meaning of its parts and the way that they are put together. But if the basic objects of computation are not themselves precise, the idea that we can *compute* the meaning of a sentence (or even a discourse) from the meanings of words is problematic from the start. Since many, perhaps most, expressions of natural language are vague, formal semantics is in deep trouble if Grim is right.

The best-known theory of vagueness that is capable of avoiding Grim’s problem is the epistemic theory [34]. 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, formal semantics can carry along merrily. However, a common objection to the epistemic theory is that it requires an implausible divorce between meaning and humans’ knowledge and use of language. In Williamson’s own words, “[a]lthough meaning may supervene on use, there is no algorithm for calculating the former from the latter”; again, “meaning may supervene on use in an unsurveyably chaotic way” [35, pp.206,209]. That is, there is no hint in the epistemic theory as to where meanings do come from. To those who view the study of language as part of (or at least closely connected to) the study of human psychology and sociology, this consequence of epistemicism tends to come across as a *reductio ad absurdum* of the theory.

Despite these criticisms, I think that the epistemic theory contains a crucial insight by treating vagueness as a result of IMPRECISE KNOWLEDGE OF LANGUAGE. That is, epistemicists locate vagueness not in the semantic theory proper, but in the relation between language users and the semantic theory. Further, since (as Williamson [34] emphasizes) uncertainty is a very general fact of the human condition, the epistemic theory has the advantage of parsimony: it allows us to reduce vagueness in language to independently motivated features of human knowledge and belief.

The theory of vagueness presented in this paper incorporates this aspect of the epistemic theory while also allowing for a close connection between meaning and use. It can be seen as a development of David Lewis’ 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” [20]. On this approach, like the epistemic theory, vagueness resides in the relation between humans and precise languages. However, it is not necessary to endorse the epistemicist’s claim that there is a single precise language that is being spoken – “English” or “Swahili”, say, or some ultra-precise variant of these. 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 (see also [19, 3]). This approach does not in itself rule out the epistemic theory;

but, I will suggest, it makes it possible to claim the advantages of epistemicism without also taking on its less palatable consequences.

The theory developed here, like Lewis’ and Williamson’s theories, holds that vagueness is not a special semantic phenomenon, but a consequence of the nature of linguistic knowledge and general principles of language use. The precise development of this claim, however, will not be in terms of “fuzzy” regions but, like Edgington [4], in terms of PROBABILISTIC linguistic representations. There is independent reason to believe that human knowledge is represented probabilistically. I argue that this perspective leads naturally to a model of interpretation as an interpreter mapping words and other utterance-types to a probability distribution over precise resolutions. This approach allows us to extend existing models of probabilistic knowledge representation and reasoning to the interpretation and representation of vague terms. In this way, we can explain how formal semantics is possible, and show not only why the conclusion of the sorites paradox is false, but also why the premises seem so compelling.

The essay is structured as follows. Section 2 describes the set-theoretic model of assertion and belief update and arguments from Stalnaker [28] and Barker [1] that this model also encompasses beliefs and assertions about language, including vague language. Section 3 argues that a probabilistic enrichment of this model is needed to account for partial belief and to avoid a sorites-like paradox for factual beliefs involving continuous sample spaces. A similar argument suggests that linguistic representations are also probabilistic. I explicate this by offering a possible-worlds model of linguistic knowledge akin to the familiar possible-worlds model, and show how this leads naturally to a picture of lexical representations as probability distributions over model-theoretic objects. Section 4 shows how the probabilistic approach resolves the sorites paradox, focusing on the differences between model-theoretic and probabilistic variants of the inductive premise and why they come apart. Section 5 points out some important features of the model relating to assertibility, negation, borderline cases, and common ground. In section 6 I argue compare the theory to alternative accounts, showing that it is preferable in terms of empirical coverage and the intuitive correctness of its predictions, and I answer objections from Stephen Schiffer and Nicholas Smith, who have claimed that partial beliefs about the applicability of vague terms do not behave like subjective probabilities. Finally, section 7 is a brief treatment of higher-order vagueness, showing that this perspective can make sense the dual role of *definitely* as an epistemic modal and a metalinguistic modifier.

2 Metalinguistic Assertion and Linguistic Knowledge

2.1 Assertion and Metalinguistic Assertion

In Stalnaker’s model, 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 [28]. Suppose you don’t know whether it is raining outside. If a reliable source tells you, “It is raining”, you will normally update your beliefs so that you no longer consider worlds in which it is not raining to be candidates for being the actual world. Taking propositions and common grounds to be sets of worlds, Stalnaker suggests treating the update operation as intersection: that is, to update the common ground with the information that p , simply intersect the current common ground with the p -worlds. Private belief update is treated similarly: if a person comes to believe that p , then their new belief state is simply the old belief state intersected with the proposition that p .

Stalnaker [28] also notes that assertions do not only convey information about the non-linguistic world (states of weather and the like). Rather, utterances may give information about how other utterances are to be interpreted. For example, suppose someone asks you, “What is an optometrist?” In this context, the reply “An optometrist is an eye doctor” serves 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 particular sequence of sounds in a certain way.

We might think that ignorance about technical vocabulary is a special case, though: perhaps, in the general case, there is a single clear-cut “current language of conversation”. Barker [1] and Stalnaker [31] argue that there is not: rather, there are typically many languages which are viable candidates for being the language of conversation, just as there are typically many worlds that are viable candidates for being the actual world. Barker’s discussion focuses on vague adjectives like “tall”. In Barker’s example, imagine that you arrive in a new town and you have no idea about the typical heights of local inhabitants. You ask a local: “What counts as ‘tall’ around here?”, and the local responds: “See John over there? John is tall.” Even if we have quite precise information about John’s height – say, we know that he is 5’11.4” – this utterance can be informative because it has the metalinguistic effect of narrowing down the range of possible interpretations of “tall”. If this utterance is true, then the local meaning of “tall” must be such that John counts as tall; so we can eliminate languages where John does not count as “tall” as candidates for being the current language of conversation.

Barker’s technical implementation of this effect is very close to the original Stalnakerian model (especially [30]). He assumes that in each world in the common ground there is a unique precise language that is the current language of conversation. Metalinguistic effects are modeled by the same update procedure as ordinary assertions (in Barker’s implementation, using the apparatus of Dynamic Semantics). So, for example, before the local’s utterance, there were worlds in the common ground where the standard for counting as “tall” in the current conversation was 6’0”, so that people who are 5’11.4” do not count as “tall”. After this utterance, these worlds are eliminated.

2.2 A Model for Metalinguistic Belief and Assertion

Barker’s model of metalinguistic assertion has one feature worth flagging here. Since each world is associated with a unique precise current language of conversation, the model effectively assumes an epistemic theory of vagueness: if we were able to discern precisely what world we are in, we would also be able to fix the precise interpretation of all vague terms. I have already indicated my discomfort with this idea: I think it extremely implausible that human linguistic practices are somehow able to determine (and without speakers’ knowledge) a single precise language that is being spoken in a given conversation. But it is not difficult to construct an alternative version of Barker’s model that does not have this commitment.

Suppose that there is a domain of worlds W and a domain of possible languages L , where each $l \in L$ is a partial function from utterance-types to model-theoretic objects. Each conversational participant $x \in X$ has a belief-set $B_x \subseteq W \times L$ – that is, beliefs are sets of world-language pairs. For any x , x ’s *factual* belief-set is

$$\{w \mid \exists l \in L : (w, l) \in B_x\},$$

and her *metalinguistic* belief-set is given by

$$\{l \mid \exists w \in W : (w, l) \in B_x\}.$$

Define the *factual common ground*, more or less standardly, as the intersection of the conversational participants' factual beliefs (and so the union of their belief-sets):

$$\bigcup_{x \in X} \{w \mid \exists l \in L : (w, l) \in B_x\}$$

The *linguistic common ground* will then be the intersection of the conversational participants' metalinguistic beliefs.

$$\bigcup_{x \in X} \{l \mid \exists w \in W : (w, l) \in B_x\}.$$

The reason for separating the two components of B_x is that they serve very different roles in conversation. The goal of inquiry is to exchange information in an effort to narrow down the domain of possible worlds [29]. The purpose of gaining metalinguistic knowledge, I take it, is to enable people to conduct inquiry more efficiently: the more we know about each others' linguistic habits, the smaller the linguistic common ground, and the more effectively non-metalinguistic assertions will be able to narrow down the factual common ground. On this model there is no "fact of the matter" about which language is being spoken that goes beyond the attempt to coordinate on a set of languages as small as possible, in order for conversational participants to exchange information as efficiently as possible. In other words, the current languages of conversation are just the languages in the linguistic common ground. There is no need for reference to some extrinsically given set of facts about which language is being spoken: choosing a (set of) common language(s) in a given conversation is a coordination game in the sense of Schelling [25] and Lewis [19].

Within this model we can treat the various types of information gain as follows. Call the informational effect of an assertion (or other event) *purely factual* if the same languages are considered possible in the prior state S and the posterior state S' : that is, if the factual common ground in S' is a proper subset of the factual common ground in S , but the linguistic common grounds at S and S' are equal. The effect of an assertion is *purely metalinguistic* if the linguistic common ground in S' is a proper subset of the linguistic common ground in S , but the factual common grounds are equal. The effect of an assertion is *mixed* if it is neither purely factual nor purely metalinguistic, i.e. if both worlds and languages are eliminated from the common ground.¹

I should add that there is no reason in principle why the proposal I will develop could not be treated as an implementation of the epistemic theory of vagueness. To construct such a variant, simply ignore the contents of this subsection and treat the set-theoretic model as Barker gives it as the base of the probabilistic enrichment I will develop. However, it should be clear that the apparent connection with the epistemic theory is not a deep commitment of the model developed here, but simply one possible interpretation.

¹ An example of a mixed assertion is if someone says "John is tall" in a context in which we are unsure both about John's height and about the interpretation of "tall". Then, for any height h such that we are sure that John is at least as tall as h , we can eliminate languages from the common ground where the standard for "tall" is greater than h . Likewise, if we are sure that the standard is at least h' , then we can eliminate worlds in which John's height is less than h' . In the probabilistic model to be developed, Bayesian update for mixed assertions is defined in similar fashion. This aspect is important because beliefs about the world and beliefs about language obviously do interact: we would not want a theory that separates them completely. See section 5.1 for the probabilistic version.

3 Probabilistic Semantic Representation

3.1 Toward a Probabilistic Model: Factual Information

The trouble with the model we have just outlined – both in the Stalnaker-Barker form and in the modified form just presented – is that it is not sufficiently rich to deal with cases where information change does not involve eliminating possibilities. You think you saw rain just now, but you consider it possible that you were mistaken. You may then increase your degree of belief in the proposition that it is raining without eliminating any worlds from the set of worlds considered possible. But if (factual) belief update is intersective, this option is excluded: you must either eliminate the worlds where it is not raining or leave them in. Intersection is not an appropriate model for this change in information, then.

Similarly, the effect of assertions is not always intersective. Consider the testimony of a witness of unknown reliability, who tells you “It is raining”. (I am supposing that the meaning of this utterance is sufficiently clear that we do not need to worry about metalinguistic effects for the moment.) You might wish to increase your degree of belief in the proposition that it is raining, without eliminating the possibility that it is not – i.e., the possibility that the witness is misinformed or lying. Again, the simple intersective model of belief update is inadequate, both to model the update and to deal with degrees of belief.

An obvious way to deal with problems of this type is to represent partial belief using an enriched model such as subjective probability. On this approach, after looking out the window you should update your beliefs about the likelihood of rain using Bayes’ rule, according to the estimated likelihood that the evidence of your senses really did indicate rain (and not, say, a sprinkler outside the window). Similarly, in the case of the unreliable witness, your belief in rain after the witness’ assertion “It is raining” will depend on your estimate of the probability that the witness is a truth-teller.

A second problem with the set-theoretic model of belief is an analogue of the sorites paradox with beliefs involving continuous sample spaces. There is a real number r such that the top of the Eiffel Tower is r kilometers away from the top of Big Ben. I know for sure that r is not less than 100, and that it is not more than 1000, but I certainly do not know what r is with any precision. However, my knowledge is even more imprecise than this characterization may suggest. Intuitively, (4) holds:

- (4) There is no real number r' such that my belief state allows for the possibility that Big Ben and the Eiffel Tower are r' kilometers apart, but excludes the possibility that they are $r' \pm \epsilon$ kilometers apart for sufficiently small ϵ .

Going through a forced march – “Are Big Ben and the Eiffel Tower 400 kilometers apart? Are they 400.01 kilometers apart?” etc. – there is no point at which I would be comfortable switching from “maybe” to “no” in a single increment of 0.01 kilometers.

In the set-theoretic model, (4) entails that there is no sharp cut-off in which worlds are considered possible: for any r' , if there are r' -worlds in the belief-set, then there are r'' -worlds in the belief-set for any $r'' \in [r' - \epsilon, r' + \epsilon]$. It follows that my belief-set contains worlds where r is any number you like, including, e.g., 1 kilometer and 1,000,000 kilometers. This contradicts the assumption that I know that r is not less than 100 or more than 1000. The paradox, in effect, is that imprecise knowledge is no knowledge at all.

Again, an obvious approach to this problem – though not the only one, to be sure – is to treat factual beliefs as probability distributions over (appropriate subsets of) W . The reason that (4) holds, on this model, is that r' is a continuous random variable, so that there are no sharp cut-offs in the likelihood that $r = r'$. Suppose,

for example, that $\text{prob}(r = r')$ is normally distributed with $\mu = 400, \sigma = 100$. Then it is much more likely that r is in $[390, 400]$ than that it is in $[190, 200]$, which is in turn much more likely than that r is in $[10, 20]$. Incidentally, since the latter has an extremely small but non-zero probability, this example technically requires a semantics for the adjective “possible” that does not equate possibility with non-zero probability. For relevant discussion see [32, 37, 38, 18].

3.2 Toward a Probabilistic Model: Metalinguistic Information

Factual beliefs involving continuous sample spaces create problems for the set-theoretic model of belief that are uncannily similar to the sorites paradox. This is not, I will suggest, accidental: standard examples of vagueness involve predicates ranging over sample spaces that are either continuous (e.g. “tall”) or involve very small increments (e.g. “heap”). I will suggest that the solution to the problem of vagueness is effectively the same as the solution to the problem in (4) suggested in the previous section.

For an example of metalinguistic belief involving a continuous sample space, we can stick with our stock example “tall”. Consider Barker’s scenario again. We know exactly how tall John is – he is 5’11.4” – and we know that he counts as “tall” in the local community. Barker’s approach to metalinguistic update encounters a problem similar to (4), given in (5):

- (5) In the context just described, there is no real number r' such that my belief state allows for the possibility that “tall” means “having height at least r' inches”, but excludes the possibility that “tall” means “having height at least $r' - \epsilon$ inches” for sufficiently small ϵ .

(We use $r' - \epsilon$ rather than $r' \pm \epsilon$ here because there clearly is an r' in the linguistic common ground such that we can be sure that “tall” does not mean “having height greater than $r' + \epsilon$ inches”: just set $r' = 5'11.4''$.)

(5) is of course a variant of the inductive premise of the sorites paradox, tailored to the set-theoretic model of metalinguistic belief. It strikes me as highly plausible: surely there is no point, as we move from 5’11.4” down to 1” at 0.01-inch increments, where I could comfortably agree that some precise interpretation of “tall” is possible, but the next interpretation is totally impossible. Add a few plausible premises and we have a full-blown sorites paradox:

- (6) a. My belief state entails that someone who is 6’9” tall counts as “tall” in any language in the common ground.
- b. My belief state entails that someone who is 1 foot tall does not count as “tall” in any language in the common ground.

(6a) and (6b) are mutually incompatible with (5). In other words, if (5) is correct, then my belief state must admit interpretations of “tall” where the standard is below 5’11.4”, even ones where “tall” means “having a height of 1 inch or more”.

This consequence is surely unpalatable: I am quite confident, for example, that someone who is one foot tall does not count as “tall” in any English-speaking community. Within this approach, our options are to reject (5) or to reject (6b). Rejecting (6b) seems to be out; but rejecting (5) would mean that there is a sharp cut-off in my metalinguistic beliefs regarding “tall” which I am unaware of (and presumably have no introspective access to). This response is no more plausible, I think, than it would be in the case of factual uncertainty.

The solution in this case should be, I think, just as in the factual case: enrich the belief space using probability measures or something with an equivalent effect. Let a *probabilistic belief space* be a triple $\langle W, L, \mu \rangle$, where W is a set of possible

worlds, L a set of possible languages, and $\mu : (W \times L) \rightarrow [0, 1]$ a function from world-language pairs to probabilities obeying the usual axioms. If we like, we can define separate probability measures for languages and worlds:

$$\mu_W(w') \stackrel{\text{def}}{=} \sum_{l \in L} [\mu(w', l)]$$

– the total probability of a world w is just the sum probability of all world-language pairs in which w occurs – and likewise

$$\mu_L(l') \stackrel{\text{def}}{=} \sum_{w \in W} [\mu(w, l')]$$

– the probability simpliciter of a language l is the sum probability of all world-language pairs in which l occurs.

I will make use of these abbreviations in what follows, but it is important to keep in mind that languages and worlds do not, in general, vary independently: facts about the world impose serious constraints on what languages are plausible candidates for receiving significant probability mass. To pick an obvious case, if we are talking about cats, no world-language pair (w', l') will receive significant probability if l' does not yield a value for the sequence of sounds “cat”, or if it assigns “cat” some value that is totally unrelated to cats. So the choice of language is obviously constrained by facts about the world, in this case facts about what noises are being made in the current conversation.

Less trivially, if someone says “Mary is tall” in a context where we are uncertain both about Mary’s height and about the interpretation of “tall”, the probabilities of worlds and languages will certainly not be affected independently by this utterance. In particular, if we are certain (i.e., probability 1) that Mary counts as “tall” in the local context and that Mary is at least h inches tall, then we can assign probability 0 to any world-language pair (w, l) in which the interpretation of “tall” in l does not include people of height h . If the probability is small but non-zero that Mary is at least h' inches tall, then we should adjust our probability for languages where “tall” receives a value of at most h' inches to some appropriately small value. (I don’t want to get involved in spelling out the complete update procedure here, but it is a straightforward application of Bayes’ rule.) In any case, the idea is that the choice of world will constrain, but not fully determine, the probabilities that can appropriately be assigned to world-language pairs.

The explanation of the modified sorites paradox for metalinguistic belief (5) within this approach will be essentially the same as the explanation of (4) given above: the value of *tall* is a continuous random variable. This interpretation of (5) says that there is no point at which the probability that “tall” means “having height at least r inches” is substantially greater than the probability that “tall” means “having height at least $r - \epsilon$ inches” for small ϵ . Nevertheless, the probability that “tall” means “having height of at least r' inches” approaches zero as r' gets smaller. The probability that “tall” applies to people who are 1 foot tall, for example, is effectively zero.

3.3 Lexical Probabilities

The notion of a possible language, though useful for our purposes, suffers from much the same psychological implausibility as the notion of a possible world: just as people do not reason about factual issues using an infinite set of discrete possibilities, they surely do not reason metalinguistically using an infinite set of discrete languages. The solution, in the factual case, is to think of beliefs as probability distributions over propositions. Probabilities of factual events can of course be defined

equivalently in either way. Similarly, we can simplify our agents' task by treating metalinguistic belief in terms of probability distributions over precise resolutions of individual words, and relate to an equivalent formulation in terms of the probability function μ (and its derivative μ_L).

An agent's belief-set determines the representation of a word or utterance-type in the following way. As usual, $\mathbf{D}_e$ is the set of possible objects whose members are $\mathbf{o}_1, \mathbf{o}_2, \dots$. For simplicity's sake we will restrict attention to model-theoretic objects of type e and $\langle e, t \rangle$, although the definition could easily be expanded to account for objects of arbitrary type.²

Define the *lexical probability function* $LP_{u,A}$ of a word u according to an agent A as a function $f_A : \mathbf{D}_e \rightarrow [0, 1]$, subject to the condition in (7):

$$(7) \quad f_A(u(\mathbf{o}_m)) = \sum_{l \in L} [\mu_L(l) : l(u)(\mathbf{o}_m) = \text{True}]$$

(μ_L should of course be understood as relativized to A 's belief state here. Note also that (7) is only appropriate for finite L ; the infinite case is not significantly different except that a bit of calculus is needed.)

In words, (7) says that the probability according to A that the word u applies to the object $\mathbf{o}_m$ is equal to the sum probability of all possible languages such that the value of u in that language, applied to $\mathbf{o}_m$, yields the value True. Since no ambiguity arises, I will use “*prob*” as a probability function for both languages and words/utterance-types from now on. The reader should understand “ $prob_A(u(\mathbf{o}_m)) = d$ ” as meaning

$$f_A(u(\mathbf{o}_m)) = d$$

or, equivalently, as

$$\sum_{l \in L} [\mu_L(l) : l(u)(\mathbf{o}_m) = \text{True}] = d$$

As an example, suppose that the lexical representation of *tall*, $LR_{tall,A}$, for a particular speaker A has the following form. A considers possible these resolutions of *tall*: 5'6", 5'7", ... 6'5" (spaced at 1" to simplify the model; the extension to continuous probability spaces is straightforward). Letting italics represent utterance-types as above and boldface indicate model-theoretic objects, we'll call these ***tall***₁, ***tall***₂, ... ***tall***₁₂. Each resolution of *tall* denotes (the characteristic function of) a set of individuals who satisfy certain conditions. So, for example, $\llbracket \mathbf{tall}_1 \rrbracket = \lambda x.x$'s height $\geq 5'6''$; $\llbracket \mathbf{tall}_2 \rrbracket = \lambda x.x$'s height $\geq 5'7''$; and so forth. The probability distribution in the bottom row of Table 1 assigns a probability in the range $[0, 1]$ to each resolution of *tall*. For example, the third column of Table 1 should be read: “The probability that the denotation of *tall* is ‘ $\lambda x.x$'s height $\geq 5'8''$ ’ is 0.03”.

² I am supposing that gradable adjectives have semantic type $\langle e, t \rangle$ (cf. [17, 1]), though nothing hinges crucially on this assumption. If gradable adjectives turn out to be of type $\langle d, et \rangle$ [33, 10] or type $\langle e, d \rangle$ [14–16], they will need special treatment. For example, von Stechow [33] and Kennedy [15] convert gradable adjectives in the positive form into properties of individuals using a silent degree morpheme. The value of this silent morpheme is given by a contextual parameter or something with equivalent effect, and is generally quite underdetermined. All of the arguments given here for the probabilistic model apply equally to the contextual determination of a value for the silent positive morpheme, I think, and we could easily extend the current approach to include probability distribution over contexts – something that is probably needed anyway. And, however gradable adjectives are best handled, the model described here is needed for vague terms like *heap* and *Mount Everest* that are not gradable adjectives, and do not show any sign of reference to degrees in their semantics.

Name	tall ₁	tall ₂	tall ₃	tall ₄	tall ₅	tall ₆	tall ₇	tall ₈	tall ₉	tall ₁₀	tall ₁₁	tall ₁₂
Threshold	5'6"	5'7"	5'8"	5'9"	5'10"	5'11"	6'0"	6'1"	6'2"	6'3"	6'4"	6'5"
$prob_A$ ($tall = \mathbf{tall}_n$)	0	.01	.03	.09	.14	.23	.23	.14	.09	.03	.01	0

Table 1. Sample lexical probability function for *tall*.

As in (7), for any word u all of whose interpretations denote sets of individuals, the probability that x is u is the sum of the probabilities of all interpretations of u of which x is an element. In the case of scalar adjectives like *tall* it is a straightforward matter to calculate the probability that an individual x is u , because all of the available interpretations of *tall* in Table 1 are upward monotonic. That is, if an individual x is a member of the set denoted by $tall_n$, and the height of an individual y is greater than that of x , then y is also in the set denoted by $tall_n$.

Thus, for any height, the probability that an individual of that height will count as *tall* is simply the sum probability of all thresholds of height less than or equal to that height, i.e. the cumulative probability. To illustrate, consider 12 individuals in this height range, spaced at 1 inch, called $x_1, x_2, \dots, x_{12}$. Using Table 1, we can calculate for each individual x_n the probability that x_n is *tall* as follows. For example,

$$prob_A(tall(x_5)) = \sum_{i=1}^5 prob(tall = \mathbf{tall}_i) = 0 + .01 + .03 + .09 + .14 = .27.$$


Table 2 gives the values of $tall(x_n)$ according to (7) and the probability distribution in Table 1 for a representative sample of individuals of various heights.

Individual	x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_8	x_{29}	x_{10}	x_{11}	x_{12}
Height	5'6"	5'7"	5'8"	5'9"	5'10"	5'11"	6'0"	6'1"	6'2"	6'3"	6'4"	6'5"
$prob_A(tall(x_n))$	0	.01	.04	.13	.27	.5	.73	.87	.96	.99	1	1

Table 2. Cumulative probability distribution corresponding to Table (1).

As Table 2 suggests, the upward monotonicity of all resolutions of *tall* in Table 1 explains the intuitive fact that the probability of an individual x_n being tall increases gradually the greater x_n 's height is. Graphically, Table 2 yields the cumulative distribution:

Since we are looking at increments of 1", this graph only approximates a smooth curve. We can increase the resolution of Tables 1 and 2 by considering more intermediate cases while maintaining the shape of the curve. The representation of *tall* (for our agent A) is a function which yields the values considered here at intervals of 1". In the limiting case in which we consider a dense scale of heights, the curve will be smooth. Essentially, I am suggesting that the distinguishing characteristic of vague predicates like *tall* is that their lexical representations are described by CONTINUOUS PROBABILITY FUNCTIONS. Non-vague terms, then, will be those whose representations are (in the relevant regions) given by discontinuous probability functions.



4 The Sorites Paradox

The original motivation for our discussion was the sorites paradox, which is restated, substituting **tall** for **heap**, in (8). As above, x_1 is 5'6", and x_{12} is 6'5". Again, boldface indicates model-theoretic objects while italics indicate utterance-types. The use of boldface **tall** in (8) makes explicit the implicit assumption in the original statement of the sorites paradox that the words in question denote unique model-theoretic objects.

- (8) a. $\neg \mathbf{tall}(x_1)$
 b. $\forall n[\neg \mathbf{tall}(x_n) \rightarrow \neg \mathbf{tall}(x_{n+1})]$
 c. $\therefore \forall n[\neg \mathbf{tall}(x_n)]$

Intuitively, premise (8a) is clearly true – someone who is 5'6" is not tall. But the conclusion that no one is tall is clearly false. Our only hope is to deny (8b); but this is strange, since this premise is intuitively very plausible.

Within the probabilistic theory of linguistic representation that I have sketched, the paradox in (8) could be restated in two ways. Suppose first that we consider the intended interpretation of **tall** as a model-theoretic object. **Tall** in itself does not represent any object in the theory proposed here: a subscript is needed to indicate which possible language $\mathbf{tall}_n$ is intended to belong to. So, if *tall* is assigned a value in l_{34} , then $l_{34}(\mathbf{tall})$ is written $\mathbf{tall}_{34}$. The paradox now appears as in (9):

- (9) a. $\neg \mathbf{tall}_{34}(x_1)$
 b. $\forall n[\neg \mathbf{tall}_{34}(x_n) \rightarrow \neg \mathbf{tall}_{34}(x_{n+1})]$
 c. $\therefore \forall n[\neg \mathbf{tall}_{34}(x_n)]$

But since possible languages are perfectly precise, the inductive premise is plainly false: whatever l_{34} is, there must be some sharp boundary between the extension of $\mathbf{tall}_{34}$ and that of $\neg(\mathbf{tall}_{34})$, and so the conclusion does not follow.

Suppose now we rewrite the paradox using words in place of model-theoretic objects. *Tall* is a word, and it is interpreted by some possible language l_n as a semantic object $\mathbf{tall}_n$. Crucially, *tall* is not in itself a model-theoretic object. If we attempt to restate the paradox using *tall* in place of **tall**, we get (10).

- (10) a. $\neg \mathbf{tall}(x_1)$
 b. $\forall n[\neg \mathbf{tall}(x_n) \rightarrow \neg \mathbf{tall}(x_{n+1})]$
 c. $\therefore \forall n[\neg \mathbf{tall}(x_n)]$

None of the clauses in (10) are well-formed within the theory I have introduced. I have occasionally spoken of “the probability that u applies to $\mathbf{o}$ (according to A)”,

but this was explicitly introduced as an abbreviatory convention. The bare claim that x_n is *tall* is meaningless: we can only say that *tall* applies to x_n with some probability d .

Suppose we rewrite the paradox using probabilities, as the present approach demands. It is plausible within a probabilistic theory (though not quite right, as we will see) to suppose that “ x_n is not tall” is expressed as “ $\text{prob}_A(\text{tall}(x_n)) = 0$ ”. If we accept this, the restatement of the sorites paradox is:

- (11) a. $\text{prob}_A(\text{tall}(x_1)) = 0$
 b. $\forall n[\text{prob}_A(\text{tall}(x_n)) = 0 \rightarrow \text{prob}_A(\text{tall}(x_{n+1})) = 0]$
 c. $\therefore \forall n[\text{prob}_A(\text{tall}(x_n)) = 0]$

(11) is logically valid, but premise (11b) is much less compelling than the original inductive premise (8b) seemed to be. 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).³

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

$$(12) \quad \neg \exists n[\text{prob}_A(\text{tall}(x_n)) = 0 \wedge \text{prob}_A(\text{tall}(x_{n+1})) = 1]$$

But (12) is not equivalent to (11b) in the current theory: 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. (12) is true of *tall* and, suitably modified, any other vague predicate, but this creates no problem: it is, if anything, just a necessary condition for a predicate’s being vague.

Delia Graff Fara has claimed that a convincing theory of the sorites, if it denies the validity of the inductive premise, must answer three separate questions (slightly modified from [5]).

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

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

- (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?

³ An influential objection to supervaluational treatments of vagueness is that, even though there is no sharp boundary between the positive extension and the negative extension of a vague predicate, the sentence “There is a sharp boundary between the positive extension and the negative extension of ϕ ” is supertrue for all predicates ϕ , whether or not they contain vague terms. Eytan Zweig points out that one could construct an analogue to the supervaluationist’s problem for my theory: the sentence “there is an n such that $\text{prob}_A(u(x_n)) = 0$ and $\text{prob}_A(u(x_{n+1})) \neq 0$ ” has probability 1 whenever, for all l_n such that $\mu_L(l_n) \neq 0$, $l_n(u)$ is defined. One possible approach is simply to say that people do not have reliable intuitions about infinitesimal differences in probability, so that the validity of this statement is not problematic. An alternative would be to deny that probabilities ever really reach 0 or 1 except for logical contradictions and validities, respectively: so, for example, the probability that a 2-foot-tall person counts as tall is infinitesimal, but not zero, and the probability that a nine-foot-tall person is tall is extremely close, but not quite equal, to 1. See the following section for more discussion of the second approach.

2. **The Epistemological Question:** If “ $\forall n[\neg \text{tall}(x_n) \rightarrow \neg \text{tall}(x_{n+1})]$ ” 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 the universal variant of the 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_n**, 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_n** but an intuition about the meaning of the word (utterance-type) *tall*.
- (b) If we replace **tall** in Fara’s formulation of the “sharp boundaries” claim by a word such as *tall*, making appropriate adjustments, the “sharp boundaries” claim is false. However, no revision of classical logic and semantics is required to explain these facts; rather, the falsity of the sharp boundaries claim follows from the fact that the word *tall* does not denote a unique object, but denotes various objects with differing probabilities. The semantic metalanguage is nevertheless classical.
2. **The epistemological question.** “ $\exists n[\neg \text{tall}(x_n) \wedge \text{tall}(x_{n+1})]$ ” is not well-formed in the present theory. If we substitute **tall_m**, as in “ $\exists n[\neg \text{tall}_m(x_n) \wedge \text{tall}_m(x_{n+1})]$ ” we can identify which n satisfies this formula given a complete specification of the language l_n or of the extension of **tall_m**. If we consider the sharp boundaries claim substituting the word *tall*, our language does not permit us to ask which n makes “ $\exists n[\neg \text{tall}(x_n) \wedge \text{tall}(x_{n+1})]$ ” true, because this sentence is not well-formed. But the probabilistic version of this formula – $\exists n[\text{prob}(\text{tall}(x_n)) = 0 \wedge \text{prob}(\text{tall}(x_{n+1})) = 1]$ – is plainly false, an intuitively correct result.
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. So it is almost always safe to assume that if an individual counts as *tall* in some context then an adjacent individual will also count as *tall* in the same context. However, informal deductions involving vague terms of natural language are not reliably truth-preserving because the terms are not associated with a unique model-theoretic interpretation.

5 Some Loose Ends

5.1 Assertibility, Joint Distributions, and Borderline Cases

When is the sentence “ x is tall” assertible? In the previous subsection we supposed that it is when $\text{prob}(\text{“}x \text{ is tall”}) = 1$. But this cannot be quite right: if the meaning of “tall” is a continuous random variable, then the cumulative probability approaches 1 in the limit as height goes to infinity, but never reaches 1 at finite height.

With respect to factual beliefs, it is commonly supposed that the norm of assertion is knowledge [36]. If this is correct, then we should expect, as a descriptive

matter, that a cooperative speaker A will typically assert things that she *thinks* she knows. On standard assumptions, A 's subjective probability for a proposition p will rarely if ever reach 1; but p may have high enough probability that A thinks that she may profitably make an assertion calculated to communicate the information that p . How high is judged “high enough” will depend on various features of the context, such as the conversational stakes and perhaps even aspects of the speaker's personality. I will use α as a placeholder for the threshold of assertibility, however this is determined in particular contexts. Cooperative speakers will assert a proposition p only if the probability of p is greater than α , and the information that p is deemed relevant, useful, etc.

The extension to vague terms is straightforward: in cases where the height of an individual x is known, and a speaker wants to decide whether to describe him as “tall”, she can simply compute whether or not $\text{prob}(\text{tall}(x)) > \alpha$; if so, “ x is tall” is assertible. We have already seen how this would be done in cases in which the individual's height is known with precision. In cases where there is both linguistic and factual uncertainty, the speaker can compute the joint probability distribution for the two uncertain variables in question – x 's height h , and the probability that someone of height h counts as “tall”. In the finite case, this is just the average of the probability that x counts as “tall” in various worlds, weighted by the probability that these worlds are the actual world.

$$\text{prob}(\text{“}x \text{ is tall”}) = \sum_{l \in L} \sum_{w \in W} [[\mu_L(l) : l(\text{tall})(x)(w) = \text{True}] \times [\mu_W(w)]]$$

That is, using the abbreviatory convention defined above,

$$\text{prob}(\text{“}x \text{ is tall”}) = \sum_{w \in W} [\text{prob}(\text{tall}(x)(w)) \times \mu_W(w)]$$

(Again, the infinite case is an straightforward extension.) So “ x is tall” is assertible just in case its probability is greater than α , taking into account both linguistic and factual uncertainty.

Since our metalanguage is classical, $\text{prob}(\text{“}x \text{ is not tall”}) = 1 - \text{prob}(\text{“}x \text{ is tall”})$. Since “ x is not tall” is assertible just in case $\text{prob}(\text{“}x \text{ is not tall”}) > \alpha$, it follows that “ x is not tall” is assertible just in case $\text{prob}(\text{“}x \text{ is tall”}) < 1 - \alpha$. A characterization of borderline cases – those for which neither “ x is tall” nor “ x is not tall” is assertible – follows immediately:

A *borderline case* of F is an individual x for which $1 - \alpha < \text{prob}(\text{“}x \text{ is } F\text{”}) < \alpha$.

5.2 Common Ground

Since introducing the probabilistic model, we have dealt exclusively with subjective probabilities, avoiding the issue of common ground. But of course a probabilistic theory of assertion, interpretation, and shared belief is needed at some point. To give a fully explicit model of this type would be a major undertaking, but a few preliminary comments may be useful.

The probabilistic model seems to require a weaker notion of common ground than Stalnaker's: in particular, common knowledge seems unattainable. We can, however, use a metric of shared belief: x and y share factual beliefs to the extent that their probability distributions over W overlap. If the ultimate goal of inquiry is to find out how the world is [29], then the ultimate goal in our model ought to be to assign probability 1 to the actual world. Of course this goal is unattainable, but we can approach it by gaining more and more information about the world, reducing uncertainty as much as possible.

In the case of linguistic beliefs, there is no “way the world is” to be discovered; however, each participant in a conversation has a belief set including a probability distribution over a set of possible languages, and they wish to share linguistic forms so that they can exchange information. One goal of metalinguistic inquiry is to make your personal μ_L overlap as much as possible with your interlocutors’ personal μ_L . However, this is not enough – after all, one way to achieve total overlap would be for all interlocutors to have a uniform distribution over L , but this distribution would be useless for communication. Another goal of metalinguistic inquiry, then, must be to assign most of the probability mass to a relatively small set of languages – i.e., to minimize linguistic uncertainty so that the comprehension process produces a manageable set of candidate interpretations. A general theory of how language users balance these needs would be useful, but is beyond the scope of the present paper. Decision-theoretic and game-theoretic models of language use and evolution seem to me to provide a promising starting point, though [19, 21, 22, 23, 11].

6 Comparison with Alternatives

6.1 Supervaluationism, Interest-Relativity, and their Kin

As discussed in section 1, the theory of vagueness outlined here treats vagueness in terms of the relation between language(s) and language users. Other theories that share this feature are, for example, Fara [5] and Barker [1]. These approaches share with the present theory the fact that the interpretation of vague terms is influenced by the discourse context, whether in Stalnakerian fashion (Barker) or by the interests of the conversational participants (Fara). However, such approaches to vagueness – and that of Kennedy [15] – have two related drawbacks. First, they treat vagueness as a special property either of particular predicates (Barker) or of the positive form (Fara and Kennedy), rather than deriving it from more general principles. Second, these theories only offer an account of vague scalar adjectives. But vagueness extends far beyond scalar adjectives; indeed, it is surprisingly hard to find examples of natural language expressions which are **not** vague. The present theory, unlike these alternatives, predicts that vagueness, not precision, should be the norm in natural language.

The theory outlined here is to some extent an elaboration of Lewis’ suggestion in [20] that vagueness is a question of language choice, rather than an issue of semantics proper. In her discussion of Lewis’ idea and Burns’ [3] defense of this approach, Keefe [13] argues that Lewis’ “pragmatic” theory is, for all practical purposes, just a restatement of the supervaluationist approach advocated by Keefe herself and by Kamp [12] and Fine [6]. It is undeniable that Lewis’ suggestion, and the elaboration I have offered, look similar in broad outline to supervaluationism. However, there are important conceptual and empirical differences. First, the theory makes no appeal to specialized semantic concepts such as “supertruth”. (Perhaps a rational agent is obliged to assign probability 1 to certain sentences, but such sentences have no special semantic status.) Second, supervaluation theory (in the relevant form) is essentially an ad hoc theory designed specially to account for vague terms. The approach advocated here has a conceptual advantage over supervaluation theory in that it does not rely on any special semantic mechanisms, either enrichments of the semantic metalanguage or otherwise unmotivated stipulations about the lexical properties of particular words. Instead, the formal treatment of linguistic representations is maximally similar to the treatment of factual beliefs, and inherits independently motivated properties of this theory.

The current theory is also preferable to supervaluationism in at least one empirical respect: supervaluation theory founders on the issue of higher-order vagueness. That is, the theory predicts that *John is tall* is vague, but *John is definitely tall* is

not vague. This is clearly incorrect: it is just as easy to construct a sorites series for *definitely tall* as it is for *tall*. Keefe [13] responds to this criticism by appealing to a vague metalanguage, so that the vagueness of *definitely tall* resides in the vagueness of the notion of an “acceptable precisification”. But as Williamson [34] notes, pushing vagueness back from the object-language to the meta-language is not a satisfactory solution to the problem. I will show in section 7 that the present theory can deal with higher-order vagueness without resorting to a vague metalanguage.

6.2 Fuzzy Logic

As noted above, a natural interpretation of the probabilistic apparatus argued for here is in terms of degrees of belief, usually treated in terms of subjective probabilities when factual beliefs are at stake. However, Schiffer [26] and Smith [27] have argued in various ways that, although we are correct to treat borderline status as involving an intermediate degree of belief that a term applies to an object, vagueness-related degrees of belief do not behave like subjective probabilities but like degrees of truth in fuzzy logic.

Schiffer’s argument is this: if the probability of p is 0.5, the probability of q is 0.5, and p and q are independent, then the probability of $p \wedge q$ is necessarily 0.25. However, Schiffer [26, p.225] claims that

when Sally believes to degree .5 that Tom is bald, thereby making him a paradigm borderline case of baldness for her, she also believes to degree .5 that he is thin, making him also for her a paradigm borderline case of thinness ... Can we expect eminently rational Sally to believe to degree .25 that Tom is bald and thin? I submit not. I submit that she’ll believe the conjunction to degree .5.

If Schiffer is correct, then the present theory is indeed on the wrong track, and fuzzy logic does better (since in fuzzy logic the degree of truth of a conjunction is the minimum of the degrees of truth of the conjuncts: $\min(0.5, 0.5) = 0.5$).

Clashes of intuition probably will not take us far, but for what it is worth, I do not share Schiffer’s intuition about this scenario. Whatever the status of Schiffer’s claim, though, his proposal has a consequence that is much more counter-intuitive than the problem it is meant to solve. If we embrace fuzzy logic for vague terms we predict that the degree to which Tom is bald AND thin is $\min(0.5, 0.5) = 0.5$, the same as the degree to which he is either bald OR thin ($\max(0.5, 0.5) = 0.5$). This is strange: surely it is more likely that he is one or the other than that he is both. In contrast, the probabilistic theory predicts, I believe correctly, that $\text{prob}(\text{Tom is bald and thin}) \leq \text{prob}(\text{Tom is bald}) \leq \text{prob}(\text{Tom is bald or thin})$.

Smith [27] gives another argument against a probabilistic theory based on a story along the following lines. Suppose your long-lost brother is coming to visit, and you know that he is either very tall or very short, but you do not know which. This situation is qualitatively different from one in which you know that your brother is a borderline case of *tall*, and yet the probabilistic theory seems to collapse the two: in both cases your subjective probability that he is tall should be roughly 0.5. Thus, Smith concludes, a theory which treats linguistic and non-linguistic uncertainty using separate machinery is preferable.

The argument is interesting, but if it is interpreted as Smith wishes, it proves too much. Using similar reasoning we could show, without bringing in issues relating to vagueness, that many partial beliefs about factual issues do not behave like subjective probabilities either. For example, suppose in a contest I will win \$1 million if I pick the winning team in a sports contest between teams A and B. Consider two cases. In the first case, I know nothing about the teams and have no basis for

choosing. In the second case I have seen these teams play each other hundreds of times and I know that they are evenly matched: each team has won precisely 50% of the games I have watched. In both cases, according to the standard theory, it is rational for me to assign probability 0.5 to the proposition that Team A will win and 0.5 to the proposition that Team B will win. However, it is clear that these two situations are qualitatively different. So, the argument would go, a theory which treats uncertainty due to ignorance and uncertainty due to statistical knowledge using completely different theoretical machinery is preferable. (Presumably fuzzy logic would not be a candidate in this case.)

I actually agree with Smith's claim that linguistic and factual uncertainty should be treated differently, at least in part: this is why we separated world and language components of belief-sets. But the problem that Smith brings up is a very general issue for the representation of uncertainty. There are numerous proposals for how to enrich probabilistic representations to deal with examples like the one just given, e.g. ranges of probabilities or upper and lower probabilities; see Halpern [9, ch.2] for an overview. However this issue is best dealt with technically, the undeniable fact that there is a qualitative difference between the two types of uncertainty in Smith's example does not show that metalinguistic uncertainty cannot be modeled using subjective probability. At most it shows that, in an enriched probabilistic model using probability ranges or the like, metalinguistic uncertainty about borderline cases will be more similar to factual uncertainty stemming from high-quality statistical information than to uncertainty stemming from having several distinct and very different options.

Another argument which favors the probabilistic approach over fuzzy logic is noted by Edgington [4, p.305]. Imagining that d and e are two objects that are both borderline cases of being red (R),

[L]et $val(Re) = 0.5$ and $val(Rd)$ be a little less than 0.5, say 0.4. What is $val(Rd \& Re)$? Here [fuzzy logic] gives a plausible answer: 0.4, the minimum of the two. But note: $val(\neg Re)$ is also 0.5 ... [and so] $val(Rd \& \neg Re)$ is also 0.4. This is immensely implausible. e is redder than d . How could it be other than completely wrong, in any circumstance, to say " d is red and e is not"? $val(Rd \& \neg Re)$ should be zero.

This is indeed implausible. The probabilistic theory gets this case right, however: because of the upward monotonicity of gradable adjectives, any resolution of red which makes " d is red" true will make " e is red" true as well, and so the probability that d is red and e is not is 0. This is, as Edgington also concludes, a strong argument for the probabilistic approach over a theory based on fuzzy logic.

Finally, we can note that, as section 7 will explore in more detail, English (like many other languages) uses the same modal adverbs to express a high degree of certainty in the truth of non-vague propositions and a high degree of certainty that a vague predicate applies to an object. We can explicate pairs like *John will clearly/definitely leave tomorrow* and *John is clearly/definitely tall* without treating these operators as ambiguous if we simply assume that they place conditions on the probability of the expressions they modify. However, if linguistic and non-linguistic uncertainty are as fundamentally different as Schiffer and Smith argue, then we must treat these operators as ambiguous, and the fact that they occur in both contexts as a semantic accident.

7 Higher-Order Vagueness and Metalinguistic Modality

The brief discussion of higher-order vagueness above points to an important fact: *definitely* is vague, no less than *tall*. I suggest treating *definitely* as we did *tall*

above, using a probability distribution over possible precise resolutions. However, *definitely* is different in that it is a modal operator. This is clear from its double life as an epistemic modal: *John is definitely coming* does not tell us about the meaning of “coming”, but about the likelihood of John doing so. In epistemic logic *John is definitely coming* is usually taken to mean something like the following: In all of the worlds which the speaker considers live options for being the actual world, John is coming. Since the model we have adopted treats linguistic and non-linguistic beliefs similarly, we might think to define an epistemic metalinguistic modal in similar terms: *John is definitely tall* means that, in each of the languages that the speaker considers live options for being the language of conversation, John is in the extension of the interpretation of *tall* in that language.

However, since *definitely* is vague, we cannot treat it as a universal quantifier over accessible worlds. The universal quantifier is not vague, and indeed the assumption that *definitely* corresponds to a strong modal in Kripke semantics was the source of the problem of higher-order vagueness in the first place. For a standard use of epistemic *definitely* we can try instead: *John is definitely coming* is true iff the probability that John is coming exceeds some high threshold. I suggest a unified treatment of epistemic and metalinguistic *definitely*: *definitely* *u* establishes a minimum threshold for the probability of *u*, cashed out either as the probability of applicability of *u* to an individual (if *u* is resolved as a type $\langle e, t \rangle$ expression) or the utterance’s truth (if *u* is resolved as an expression of type *t*).

Definitely takes a constituent as an argument and returns a function which quantifies over the meaning of that constituent in all relevant possible languages. Since *definitely* is vague, there will be a range of resolutions available with different thresholds, just as there were for *tall*. One possible resolution of *definitely tall* – call it **definitely₆ tall₆** – is given in (13):

$$(13) \quad [\mathbf{definitely}_6 \mathbf{tall}_6](x) \text{ is true iff } \text{prob}_A(\text{tall}(x_n)) \geq 0.97$$

Definitely tall as interpreted by a language l_7 is true of *x* iff the sum of the probabilities of all languages that make *tall* true of *x* is greater than some threshold, here 97%.

On this interpretation, *definitely* is effectively a special type of epistemic modal. (Note that, just as the actual world determines the truth-value of an epistemic modal only indirectly, so the resolution of *tall* in l_n plays no privileged role in calculating the meaning of *definitely tall* in l_n .) We may suppose that the metalinguistic use of *definitely* is equivalent to the ordinary epistemic use, except that in the latter case the linguistic facts are held constant while the non-linguistic facts vary, while in the metalinguistic case the non-linguistic facts are held constant while the linguistic facts vary. Thus, the simplest extension of the theory proposed here predicts that *definitely* should be vague. It does seem that *definitely* is at least gradable, in that *John will very definitely come* is acceptable, while *very* is unacceptable with non-gradable adjectives and adverbs like *geological* [15]. A detailed comparison of metalinguistic and epistemic *definitely* will have to wait for future work; but see Sauerland & Stateva [24] for a detailed consideration of some closely related issues.

Imagine that the common ground contains seven resolutions of *definitely* with the probability distribution in Table 3. Applying (13) (with appropriate replacements) and the probabilities in Table 3 to the probability distribution for *tall* in Table 1 we get Table 4.

In general, *definitely u* will always have probability less than or equal to the probability of *u* alone, which is the correct result. Note also that iterated *definitely* is also unproblematic, and simply lowers the probability further as compared to *u* alone.⁴

⁴ A reviewer asks: Can this approach account for the possibility that *x* is neither definitely definitely tall nor definitely not definitely tall? It can. Let $\alpha = .8$ and *definitely* and

Name	def ₁	def ₂	def ₃	def ₄	def ₅	def ₆	def ₇
Threshold	≥ 0.82	≥ 0.85	≥ 0.88	≥ 0.91	≥ 0.94	≥ 0.97	$= 1$
$prob_A(\text{definitely} = \text{definitely}_n)$	0	.05	.1	.15	.2	.3	.2

Table 3. Sample probability distribution for *definitely*.

Individual	x ₁	x ₂	x ₃	x ₄	x ₅	x ₆	x ₇	x ₈	x ₉	x ₁₀	x ₁₁	x ₁₂
Height	5'6"	5'7"	5'8"	5'9"	5'10"	5'11"	6'0"	6'1"	6'2"	6'3"	6'4"	6'5"
$prob_A(\text{tall}(x_n))$	0	.01	.04	.13	.27	.5	.73	.87	.96	.99	.1	1
$prob_A([\text{definitely tall}](x_n))$	0	0	0	0	0	0	0	.05	.5	.8	1	1

Table 4. Cumulative probability distribution for *definitely tall*.

To sum up, the theory advocated here offers the promise of an explanation of higher-order vagueness with only a slight modification of existing theories of modality. This is a considerable advantage over supervaluationist treatments, which must assume that the semantic metalanguage itself is vague (as in Keefe 2000). Note also that, if this approach is viable, it constitutes a counter-example to Grim's [8] claim, quoted above, that vague terms cannot not be modeled accurately in a precise metalanguage.

8 Conclusion

The theory of vagueness described here has an important advantage over many competitors: it stipulates no special semantic apparatus for vague terms, as does supervaluation theory. Nor does it rely crucially on the claim that the meanings of words float free of speakers' knowledge of language, as does epistemicism. Rather, the probabilistic account relies on general and independently motivated properties of language use and human cognition. This account also yields an account of the sorites paradox which is explanatory with respect to Fara's three questions. Empirical advantages of the present approach over competing theories include its ability to account for vagueness outside of the realm of gradable adjectives, intuitively correct results with conjunctions and disjunctions of sentences containing vague terms, and avoidance of problems with higher-order vagueness.⁵

tall be resolved as in the text. x_{11} and x_{12} count as *definitely tall* because *tall* applies to them with probability greater than α . By reapplying Table 3 to the third line of Table 4 we see that x_{12} counts as *definitely definitely tall*, but x_{11} does not. x_9 (and all shorter individuals) count as *not definitely tall*, since $prob(\text{not definitely tall}(x_9)) = 1 - prob(\text{definitely tall}(x_9)) = 1 - .05 = .95$. However, x_9 does not count as *definitely not definitely tall* ($prob = 0.5$); only x_8 and shorter enjoy this privilege. So x_9 is neither definitely definitely tall nor definitely not definitely tall.

⁵ Thanks to Chris Barker, Philippe Schlenker, Gregory Guy, Chris Potts, Paul Egré, Stephanie Holt, and a reviewer for the current volume for helpful comments and discussion. Thanks also to Joey Frazee and David Beaver, whose very interesting paper [7] reached me too late to affect the content, but has important overlap with the current work.

References

1. Chris Barker. The dynamics of vagueness. *Linguistics and Philosophy*, 25(1):1–36, 2002.
2. Anton Benz, Gerhard Jäger, and Robert Van Rooij. *Game theory and Pragmatics*. Palgrave Macmillan, 2005.
3. L.C. Burns. *Vagueness: An Investigation into Natural Languages and the Sorites Paradox*. Kluwer, 1991.
4. D. Edgington. Vagueness by degrees. In R. Keefe and P. Smith, editors, *Vagueness: A Reader*, pages 294–316. MIT Press, 1997.
5. Delia Graff Fara. Shifting sands: An interest-relative theory of vagueness. *Philosophical Topics*, 20:45–81, 2000.
6. K. Fine. Vagueness, truth and logic. *Synthese*, 30(3):265–300, 1975.
7. Joey Frazee and David Beaver. Vagueness is rational under uncertainty. To appear in *Proceedings of the 17th Amsterdam Colloquium*, 2010.
8. Patrick Grim. The buried quantifier: An account of vagueness and the sorites. *Analysis*, 65(2):95, 2005.
9. Joseph Halpern. *Reasoning about Uncertainty*. MIT Press, 2003.
10. Irene Heim. Degree operators and scope. In Fery and Sternefeld, editors, *Audiatur Vox Sapientiae: A Festschrift for Arnim von Stechow*. Berlin: Akademie Verlag, 2001.
11. Gerhard Jäger and Robert van Rooij. Language structure: Psychological and social constraints. *Synthese*, 159(1):99–130, 2007.
12. Hans Kamp. Two theories about adjectives. In E. Keenan, editor, *Formal semantics of natural language*, pages 123–155. Cambridge University Press, 1975.
13. R. Keefe. *Theories of Vagueness*. Cambridge University Press, 2000.
14. Chris Kennedy. *Projecting the adjective: The syntax and semantics of gradability and comparison*. PhD thesis, U.C., Santa Cruz, 1997.
15. Chris Kennedy. Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*, 30(1):1–45, 2007.
16. Chris Kennedy and Louise McNally. Scale structure, degree modification, and the semantics of gradable predicates. *Language*, 81(2):345–381, 2005.
17. Ewan Klein. A semantics for positive and comparative adjectives. *Linguistics and Philosophy*, 4(1):1–45, 1980.
18. Daniel Lassiter. Gradable Epistemic Modals, Probability, and Scale Structure. To appear in *Proceedings of Semantics and Linguistic Theory XX*, 2010.
19. D. Lewis. *Convention: A Philosophical Study*. Harvard University Press, 1969.
20. D. Lewis. General semantics. *Synthese*, 22(1):18–67, 1970.
21. A. Merin. Information, relevance, and social decisionmaking: Some principles and results of decision-theoretic semantics. In L.S. Moss, J. Ginzburg, and M. de Rijke, editors, *Logic, Language, and Computation*, volume 2, pages 179–221. CSLI Publications, 1999.
22. Prashant Parikh. *The Use of Language*. CSLI Publications, 2001.
23. A.V. Pietarinen. *Game theory and linguistic meaning*. Elsevier, 2007.
24. U. Sauerland and P. Stateva. Scalar vs. epistemic vagueness: Evidence from approximators. In M. Gibson and T. Friedman, editors, *Proceedings of Semantics and Linguistic Theory XVII*. CLC Publications, Cornell University, 2007.
25. T. Schelling. *The Strategy of Conflict*. Harvard University Press, 1960.
26. S. Schiffer. Vagueness and partial belief. *Philosophical Issues 10: Skepticism*, pages 220–257, 2000.
27. N. Smith. Degree of belief is expected truth value. In R. Dietz and S. Moruzzi, editors, *Cuts and Clouds: Vagueness, Its Nature and Its Logic*. Oxford University Press, 2010.
28. R. Stalnaker. Assertion. In P. Cole, editor, *Syntax and Semantics 9: Pragmatics*. Academic Press, 1978.
29. R. Stalnaker. *Inquiry*. MIT Press, 1984.
30. R. Stalnaker. On the representation of context. *Journal of Logic, Language and Information*, 7(1):3–19, 1998.
31. R. Stalnaker. Assertion revisited: On the interpretation of two-dimensional modal semantics. *Philosophical Studies*, 118(1):299–322, 2004.
32. Eric Swanson. *Interactions With Context*. PhD thesis, MIT, 2006.

- 33. A. von Stechow. Comparing semantic theories of comparison. *Journal of Semantics*, 3(1):1–77, 1984.
- 34. T. Williamson. *Vagueness*. Routledge, 1996.
- 35. T. Williamson. Schiffer on the epistemic theory of vagueness. *Noûs*, 33:505–517, 1999.
- 36. T. Williamson. *Knowledge and its Limits*. Oxford University Press, 200.
- 37. Seth Yalcin. Epistemic modals. *Mind*, 116(464):983–1026, 2007.
- 38. Seth Yalcin. Probability Operators. To appear in *Philosophy Compass*, 2010.