GRADED INDIVIDUALS, VAGUE DISCOURSE-ENTITIES

Galit W. Sassoon
ILLC, University of Amsterdam & Ben Gurion University of the Negev
galitadar@gmail.com
Office: Science Park 904 (room C3-136K), 1098 XH, Amsterdam
PO box 94216, 1090 GE Amsterdam, The Netherlands, +31 (0)20 5258249
Home: Haverstede 19, 1112 HR, Diemen, The Netherlands, +31-(0)20-7552273

1 Introduction

1.1 Structure and goals of this paper

For the purpose of the analysis of natural language expressions, we have to adopt semantic models representing:

(i) Vagueness pertaining to the sets of instances denoted by vague predicates, and
(ii) Vagueness pertaining to the degrees that gradable predicates in degree constructions assign to individuals (as will shortly be demonstrated).

Let us call such models *vagueness models with degrees*. The rest of Part 1 presents two types of vagueness models – a Stalnakerian and a Supervaluationist model – and the way vague degree functions can be incorporated into them.

Based on this background, this paper aims to examine from close the nature of the inhabitants of vagueness models with degrees, namely the individuals that we represent as elements of the domain of discourse D_x in vagueness models with degrees of the sort described in Part 1. This examination sheds new light on two central problems, much discussed in the philosophical literature today. However, I address these problems from the perspective of a linguist, rather than a philosopher, i.e. my main interest is in the ontology presupposed by natural languages, the one whose use yields correct interpretations for natural language expressions.

The first problem, discussed in Section 2.1, is the much debated question regarding whether individuals in different worlds can ever be identified. In Section 2.2, I present a distinction between two types of circumstances. In one type of circumstances, cross world identity seems implausible, in line with Lewis's (1986) view. However, in the second type of circumstances, I claim, cross world identity does seem intuitively plausible.

The second problem, discussed in section 3.1, is the much debated question regarding whether vagueness can be identified with ignorance. While I fail to see how this problem can be solved, and do not attempt to solve it, in Section 3.2, I support the following two claims. On the one hand, the very same formal models can be used to represent both phenomena, in line with, e.g. Williamson's (1994) view that vagueness *is* ignorance. On the other hand, the proposal set out in Part 2 regarding the representation of individuals in vagueness models illuminates a sharp difference between phenomena we tend to call 'vagueness' and phenomena we tend to call 'ignorance', i.e., the intuitive distinction between the two is captured.

1.1 Models for the interpretation of vague predicates

Let us call the common knowledge of some community of speakers out of the blue an *actual context* (Stalnaker 1978).

When we look at natural languages, we see that for some predicates, like *even number*, it is determined for each individual in each actual context whether it falls under the predicate or not. There is no denotation gap containing individuals that one does not know if they fall under the predicate or not. These predicates are called *non-vague,* or *sharp*. Other predicates do have a gap. For example, in a certain domain of entities (say, the students in a certain class), we are usually able to positively classify some entities as clearly *tall*, and some other entities as clearly *not-tall*. However, for yet another set of entities, whose height falls in between that of the clearly tall and the clearly short entities in the domain, we are often not able to decide whether they are tall or not (i.e. whether in the given context c they fall within the positive denotation of *tall*, $[[tall]]_c$ or the negative denotation, $[[not\ tall]]_c$). These predicates are called *vague*. Regarding nouns like *bird* or *apple*, linguists usually assume that, for all practical purposes, it is good enough to consider them as sharp, but clearly, they are only "by and large" sharp, i.e. sometimes they do admit borderline cases (Murphy, 2002).

In attempt to capture the absence of sharp boundaries for vague predicates such as *tall, wide, heavy, tasty* and *happy,* traditionally, linguists use *Supervaluationist models* for their interpretation (e.g., van Fraassen 1969; Fine, 1975; Kamp, 1975; Klein 1980; Kamp and Partee, 1995; Veltman, 1984; Landman 1991; van Rooij, to appear, etc). Briefly, a supervaluationist model M_C consists of a set of contexts C, a domain D and a 'completion' relation between contexts. All and only completions t in $T \subseteq C$ are classical valuations: Every statement is either true or false in each t. Let $T_c \subseteq T$ be the set of completions t consistent with the information in a partial context $c \in C$:

(1) a. A statement S is *true* in c iff S is true in every completion t of c ($t \in T_c$);
b. A statement S is *false* in c iff S is false in every $t \in T_c$, and
c. The truth value of S is *undetermined* otherwise.

Another prime example of a model representing partial information (whether due to vagueness or due to epistemic ignorance) is a *Stalnakerian model* (cf., Stalnaker, 1978). The similarity in the way the interpretation of vague predicates is represented in Supervaluationist and Stalnakerian models is straightforward, as illustrated in (2a) versus (2b), respectively. For any vague predicate P in a context of partial information c:

(2) a. Supervaluationist models (van Fraassen, 1969; Kamp, 1975):
(i) The partially determined *positive denotation* $[[P]]_c$ equals the intersection of P's totally determined positive denotations, $[[P]]_t$, in c's completions $t \in T_c$ (the total contexts t consistent with c): $[[P]]_c = \cap\{[[P]]_t \mid t \in T_c\}$
(ii) The partially determined *negative denotation* $[[\neg P]]_c$ equals the intersection of P's totally determined negative denotations, $D - [[P]]_t$, in c's completions $t \in T_c$: $[[\neg P]]_c = \cap\{D - [[P]]_t \mid t \in T_c\}$
b. Stalnakerian Models (Stalnaker, 1978):
(i) The partially determined *positive denotation* $[[P]]^+_c$ equals the intersection of P's totally determined positive denotations, $[[P]]_w$, in the worlds consistent with the information in c, $w \in W_c$: $[[P]]_c = \cap\{[[P]]_w \mid w \in W_c\}$

(ii) The partially determined *negative denotation* $[[\neg P]]_c$ equals the intersection of P's totally determined negative denotations, $D - [[P]]_w$, in the worlds consistent with c: $[[\neg P]]_c = \cap\{D - [[P]]_w \mid w \in W_c\}$

Thus, while linguists typically use supervaluationist models for the representation of vagueness, and Stalnakerian models for the representation of epistemic modality, for most linguistic purposes, the predictions that can be derived from the two model types are very much alike, despite the fact that they differ with respect to the indices of evaluation.[1] In this paper, I explore the representation of individuals in vagueness models of these two sorts, i.e. in demonstrating my points I sometimes use worlds and sometimes (mainly) supervaluations (completions t of actual contexts c). I explicitly relate to the difference between the two model types when (and only when) it is particularly relevant to the points under discussion.

1.2 Models for the interpretation of Gradable predicates

The term *gradable* is used in linguistics to refer to predicates like *tall, wide, heavy, tasty healthy,* and *clever*, which are distinguished, first and foremost, by their ability to felicitously combine with comparative morphemes (as in *Dan is more / less clever than Sam is*), equatives (as in *Dan is as tall as Sam is / Dan and Sam are equally tall*), superlatives (as in *Dan is the cleverest / the least clever*), and degree modifiers (as in *Dan is very / fairly / enough / too clever*). A few adjectives (for instance, *extinct, even, married* and *nuclear*) and most nouns are classified as *non-gradable*, because they cannot felicitously combine with the given comparison and degree morphemes.

The semantic distinction underling the characteristics of gradable predicates is that entities can possess the properties that they denote to different extents. Intuitively, entities are judged to be instances of these predicates iff the extent to which they satisfy the relevant gradable property, usually called the predicate's *ordering dimension*, is within the norm, that is, iff they reach the *standard for membership* of that predicate.

Thus, in attempt to capture the fact that gradable predicates apply to entities to different extents, traditionally, linguists use degree functions for their interpretation. The use of degrees goes back to Russel, 1905, and a prime example is von Stechow (1984a,b). In each total context t a predicate P is associated with a *degree function*, f(P,t), i.e. a mapping of entities x from the domain of discourse, D_x, to degrees. Thus, e.g., *tasty* holds true of an object x in t iff the value f(tasty,t)(x) exceeds the standard of membership of *tasty* in t, standard(tasty,t). According to a prominent view, degrees are members of a linearly ordered dense set, isomorphic to the real numbers (see among many others: Klein, 1991; Kennedy, 1999; Schwarzschild, 2005; for recent studies independently motivating this approach see Fox and Hackl, 2006; Fox, 2007; Nouwen, 2008; Sassoon, 2007, 2009a,b).

Significantly, vagueness is not constrained to denotations. It may also affect the degrees assigned to entities by a given predicate. Vagueness pertaining to degrees may result due to, for instance, indeterminacy concerning the predicate's contextually relevant ordering dimension(s), or the relevant measurement method for a given dimension. Thus, degree functions often vary between contexts, e.g., the degree functions associated with *similar to Dan* may represent similarity relative to different dimensions, the degree functions associated with *tall* or *warm* may

[1] Total contexts may include more information than worlds. Thus, each world in the set W_c (that for all that we know in c may still be the actual world) corresponds to at least one completion t of c (see more on this point in 2.2).

represent different measurement techniques (different rulers or thermometers), the degree functions associated with *tasty* in different completions t of a context c may represent different tastes, etc. Thus, if we opt to give a correct representation of natural language semantic interpretations, our models should allow at least that much vagueness:

(3) Predicates' degree functions may vary between completions of actual contexts. Therefore, the value a predicate P assigns to an entity x is undetermined in a partial context c, unless it is (conventionally fixed, and hence) index-invariant: $\forall t_1,t_2 \in T_c$, $f(P,t_1)(x) = f(P,t_2)(x)$.

On this paper, then, I use the term *vagueness models* to refer to models where denotations as well as degree functions vary between indices (total completions of a given context c). Having set up this background, let us turn right now to the first part of our program, namely, the implications of a representation of vagueness about degrees on the plausibility of cross-world identity.

2. The inhabitants of vagueness models with degrees

2.1 Cross world counter-part relations

While standard semantic models represent the fact that vague predicates like *tall* have denotation gaps by assuming that their denotations vary across indices, standard semantic models in linguistics assume that proper names are rigid designators. In all completion t in T_c of a given actual context c, a proper name like *Dan* is linked to one and the same individual, say – d_1.

On the one hand, this common practice is appealing to linguists given the (perhaps wrong, but highly intuitive) view that individuals may be identified across different indices. For example, we can imagine ourselves being one centimeter taller than we are, i.e. ourselves inhabiting also an alternative reality, in which we are taller. In addition, the assumption that individuals leave a double life in two different indices appears crucial for representations *de re*, given the standard assumption that, for instance, *The dean could have been taller than she is* is true de re in an alternative of reality t iff some individual d is the actual dean in t and d's height in some accessible alternative t' is bigger than d's actual height in t. We see that the very same individual, d, is presumed to be part of both t and t'.

On the other hand, Lewis (1986) extensively argues against the plausibility of this view. Rather, for Lewis, a proper name like *Dan* denotes different individuals in different alternatives of reality, with these individuals being bound to each other by the counter-part relation. The intuition is that if in one alternative of reality (e.g., total context t) the proper name *Dan* denotes an apple and in another alternative it denotes a chair, or if in one it denotes an object that is 1.87 meters tall, and in another it denotes an object that is 1.86 meters tall, it does not make sense to say that this is the same individual. I agree with this intuition; I really do not see how we could consider them "the same". If in a given partial context c, we do not know whether a proper name like *Dan* denotes an apple or a chair, to me this means that we do not yet know to which individual in the actual total context this name refers. I represent this by saying that the extension of *Dan* is empty in c. In two total contexts in T_c, the extension of *Dan* consists of two different individuals, even if the two individuals are bound by a counter-part relation.

Certainly, Lewis views individuals and worlds as mereological sums, thus not objecting to the idea that two worlds (or completions, for that matter) share a part in common, in just the same

way as Siamese twins may share a hand. But, of course, the hand two twins have in common cannot have a different number of fingers for each twin. Thus, what Lewis does find problematic is the idea that the common part of two completions (e.g., our given object d) has different properties (heights) in the different alternatives, t and t'. Similarly, claims Lewis, something may persist by having different temporal stages d_{s1} and d_{s2} (and so it may have different, e.g., heights, in different time points, s_1 and s_2), but nothing can be wholly present at more than one time. Thus, no two different stages of an individual are one and the same, $d_{s1} \neq d_{s2}$.

In order to abstract away from the problem of change, in this paper I focus on stages of individuals in a given time point (say s), symbolizing them simply as d_1, d_2, etc., omitting the time index s. Next, I show that a close examination of the nature of the inhabitants of vagueness models makes salient an important intuitive distinction concerning circumstances under which we tend to judge individuals (focusing on individual stages per a given time point s) as "the same" across different alternatives of reality, as opposed to circumstances under which such a judgment is more difficult to justify.

2.2 Cross world identity

Significantly, Lewis argues that trans-world identity is to be rejected as a general theory of representation de re, as it falls short of meeting our needs for representations de re, *even if a couple of special cases of overlap between worlds may be tolerated*. For example, Lewis points out that if universals exist[2], then we needs to qualify his denial of trans-world identity (Lewis 1986: 69), but this qualification is innocent, as it does not pertain to any particulars. No two worlds have any one particular in common.

Importantly, Lewis distinguishes between *intrinsic* and *extrinsic* (*'relational'*) properties; the former, e.g. height, shape and number of fingers, being a matter of the way an individual d itself *is* in an alternative of reality t, while the later, e.g. *tall*, being a matter of d's *relations to others* (e.g., those representing the cutoff point between *tall* and *not-tall*) in t. In this section, I explore the consequences of the assumption that standard-dependent properties such as the one denoted by *tall* are extrinsic (or relational).

I said that if the name *Dan* denotes in one total context an object that is 1.87 meters tall and in another an object that is 1.86 meters tall, intuitively, it does not make sense to say that this is the same individual. I would now like to point out that there is a sharp contrast between that and the following case. Let us suppose that in two total contexts the proper name *Dan* denotes an object that is 1.87 meters tall, but in one total context this object is considered *tall* and in the other it is considered *not tall*, and the references of *Dan* in the two total contexts are identical in all other respects. The intuition now is that we definitely *can* say that this is the same individual. It is only our interpretation of the word *tall* that has changed. I would like to represent this intuition and account for the contrast between the two cases.

The crucial difference between the two cases seems to lie in whether 'real' properties are being compared across total contexts, rather than only the ways in which we humans categorize an individual's properties. I therefore assume the following. Two individuals in two total contexts are the same iff they have the same length, width, color, intelligence, and so on and so forth. Furthermore, the two individuals are the same in these two total contexts *regardless of whether or not* their length, width, color, intelligence and so on are sufficient for them to count as *long,*

[2] Universals are considered to be things that are wholly present as non-spatiotemporal parts in each of the things that instantiate some perfectly natural property (Lewis, 1986: 67).

wide, red, intelligent, etc. in these total contexts. In short, an object should be identified with its 'real' properties, but not with its 'linguistic' or 'conceptual' properties (so to speak).

Thus, I propose that individuals are distinguished by their *property values* (the extents to which they satisfy properties such as those denoted by gradable adjectives). I take individuals to be real entities, identified with their 'real' properties. So it is invariably determined for all individuals in D, e.g., what their heights are (and, in turn, these heights systematically constrain the values of *tall*'s degree function in every t). However, when we use proper names, we do not know exactly which individuals in D they refer to, since we do not know all of their property values.

(4) An index invariant Domain D_x of Graded Individuals:
Individuals are distinguished by their *property values* (the extents to which they satisfy properties such as those denoted by gradable adjectives). It is invariably determined for each individual in D what its property values are (its values are index invariants).

Recall that, with Lewis, I do not object to the idea that two alternatives of reality share a part in common. What is problematic for me is the idea that the common part (e.g., a given object d) has different *intrinsic* properties in the different alternatives, t and t' (at time point s), e.g., if d_1 and d_2 have different heights in t and t', I cannot consider them one and the same individual. But if two individuals have the same intrinsic properties – height, width, weight, health (in all the health respects – blood pressure, pulse, fever, etc.), happiness (in all the respects), and so one and so fourth – I consider them one and the same. Thus, if two completions t and t' may differ only with respects to the standards of vague predicates, they may share all their parts (have all the individuals in common), though the status of each individual in extrinsic, relational properties – such as *tall, popular, happy, healthy with respect to blood pressure, far away from Rome* and so on and so fourth – may vary between t and t'.

In fact, Lewis's main objection for cross world identity regards 'accidental' intrinsic properties, such as, e.g., height and shape (Lewis 1986, section 4.2.), i.e. intrinsic properties that tend to vary across alternatives of reality and time points. However, to the extent that these can be kept equal across two alternatives (or two time points), the two alternatives (or time points) may overlap (share parts).

In regard to this point, notice that Graff (2000: 57-59) argues convincingly in support of the view that the standard of vague predicates (i.e., of relational properties) may vary given a fixed comparison class, namely, when the set of entities whose height is under discussion remains fixed in the context of use, and so are their heights. This may happen if, despite the fixed comparison class, speakers' interests and goals vary, and the instances seen as typical or paradigmatic of the given predicate, as well as the instances seen as silently similar to those paradigmatic cases, vary accordingly. Thus, according to Graff (2000), the variation in standards is constrained, *but by no means fully determined,* by (i) what we may call in the context clear-cases of positive instances; (ii) the standards of antonyms of the given predicate in the given context, and (iii) the similarity constraint – If two things are saliently similar, then the standard cannot be such that one exceeds it, while the other does not. Importantly, for Graff (2000), equal differences in height may be seen as salient given one goal and not salient given another goal. Hence, we have some leeway in our standards of use for vague predicates given a certain comparison class with a given set of heights.

On the same line, I argue that as long as all accidental intrinsic properties of individuals (e.g., their heights, shapes, etc.) are kept equal across two alternatives of reality, the two alternatives may overlap (share individuals in common), despite their being distinct with respect to, e.g., the property *tall* denotes in each one of them. Alternatives of reality differing in this way (e.g., only with respect to the cutoff point between *tall* and *not-tall*) represent different speakers' goals. While they do not differ with regard to the heights under discussion, they do differ with regard to the heights forming the clear cases of *tall*, and the pairs of heights saliently similar/ dissimilar with respect to *tall*.

I have to add a small comment now about a certain difference between worlds and contexts of total-information (completions), which may be of relevance to our point. The cutoff points of vague predicates, you may claim, are not at all part of the actual world, or of worlds in general. This suggests that all the worlds in which all the entities' accidental intrinsic properties (e.g., height and shape) are identical collapse into one and the same world. If you stand behind that claim, you have to admit that completions of an actual contexts c (supervaluations) must be something richer than a world, in that in addition to the facts prevailing in a given world, they also contain information about possible cutoff points (see a supervaluationist argument to that effect in 1.2 below). That cutoff points, and hence 'rich' indices (each standing for, at least, a world plus a set of cutoff points for vague predicates) play a role in interpretation is assumed in the linguistic literature across the board. Thus, we are still in a position to claim that different indices of evaluations (completions t of contexts c in vagueness models) have overlapping parts (individuals in common). The same holds true of worlds with cutoff points for vague predicates.

Let us call my proposal that individuals are distinguished by (and only by) their values in intrinsic properties – the *relational cross-completion identity* thesis. (If you believe the standards of vague predicates are part of the world, you may also call it the *relational cross-world identity* thesis). In this section I claimed that one of its advantages is that it captures the intuitive distinction between completions in which the referents of proper names differ in height (i.e., they necessarily form different individuals) versus completions in which they do not differ in height or any other intrinsic property value (and so they do form the same individual). This intuitive distinction is blurred otherwise.

On the one hand, we regard most of the intrinsic-property values of the referents of proper names as accidental (i.e., index-variant). Thus, no proper name is a rigid designator. In particular, proper names denote different individuals in different completions of actual contexts.

On the other hand, the referent of a given proper name is intuitively considered one and the same across completions whereby its intrinsic properties are invariant. Therefore, to the extent that such completions exist, they overlap. Some individual may be part of all of them. As completions (and perhaps also worlds) may differ only with regard to the standards of vague predicates, cross-completion (and perhaps cross-world) identity of individuals is sometimes possible.

Next, I claim that, my proposal has an additional advantage, namely, it allows a natural representation of two different sources of truth value gaps – completion-variance in the interpretation of proper names versus completion-variance in the interpretation (standards of membership) of gradable predicates. While the former gives rise to what we tend to call 'ignorance', the latter gives rise to what we call 'vagueness'.

3. Two sources for truth value gaps

3.1 Vagueness versus Ignorance

The adjectives in the examples in (5) (e.g., *tall*, *intelligent* and *happy*) are vague, i.e. competent speakers of English are unable to determine whether they apply to certain entities or not.

(5) a. Dan is tall
b. Sam is intelligent
c. Mary is happy

As a result, even speakers who are familiar with all the relevant facts (e.g., Dan's height, Sam's IQ, etc.), may not be able to determine the truth values of statements such (5a-c). The most common view explains these truth-value gaps by virtue of the intuition that in some cases *there is no fact of the matter as to whether these adjectives apply or not*, i.e. these truth value gaps are not due to ignorance. The given statements lack a truth value (they are neither true nor false), so there is nothing there to know or to be ignorant of (Kamp, 1975; Kamp and Partee, 1995). Entities fail to be in the positive or negative extension of, for instance, *apple*, when they "qualify neither clearly as apples, nor clearly as non-apples" (Kamp and Partee 1995: 148), not because it is not *known* whether they are apples or not.

Conversely, the predicates in statements such as those in (6) (e.g., *two meters tall*, *taller than Sam is*, and *twice as tall as Sam*) are usually thought not to be vague. Rather, they are thought to have precise application conditions so that truth values exist for statements like (6a-c). Speakers who are not familiar with all the relevant facts (e.g., Dan and Sam's heights), and hence do not know these truth values, are therefore in a state of ignorance.

(6) a. Dan is two meters tall
b. Dan is taller than Sam is
c. Dan is twice as tall as Sam

Yet, the same formal means, e.g., Supervaluationist models (van Fraassen 1969; Fine, 1975; Kamp, 1975; Kamp and Partee, 1995; Veltman, 1984; Landman 1991, etc.) serve both for the representation of *vagueness* (e.g., the analysis of vague adjectives in Kamp, 1975 and the analysis of vague nouns in Kamp and Partee, 1995) and for the representation of *epistemic states* (e.g., Veltman's 1984 analysis of indicative conditionals and epistemic modal like *must*). Thus, the two phenomena, *vagueness* (or partial information) and *ignorance* (or partial knowledge), are often not formally distinguished.

The intuitive distinction between them is also a matter of controversy. According to, for example, Sorensen (1988), Williamson (1994) and Graff (2000), vagueness is a purely epistemic phenomenon, consisting of the absence of knowledge and of the absence of a possibility to gain knowledge concerning denotation membership and truth values, despite the existence of a true answer. For example, Williamson advances the view that the cutoff between tall and non-tall entities is determined, but unknown ('unknowable'). For each and every entity, he claims, it is determined whether it is tall or not (*tall* does not have a denotation gap), but, still, for some entities we cannot tell whether as a matter of fact they are tall or not, and for that reason we

consider them as gap members (Krifka, 2007, is an example of the influence of the epistemic view in theoretical linguistics).

I do not attempt to give arguments for or against an epistemic approach to vagueness.[3] I personally question whether the debate surrounding the epistemic approach can be settled at all. I therefore leave this debate aside. But I do attempt to address the underlying linguistic facts, i.e. the intuitions that speakers have about truth-value gaps of the sort related to the statements in (5) as opposed to their intuitions about truth-value gaps related to the statements in (6). Speakers tend to regard these two sorts of gaps as stemming from different sources. It is this fact that I aim to explain. By considering the ontology assumed in linguistic theories of gradability and comparison (degree ontology), I attempt to show that a unified (supervaluationist) model can be used to model both phenomena that are typically classified as vagueness (to be discussed in the following sections), and phenomena that are typically not classified as vagueness, but as epistemic (e.g., absence of knowledge concerning truth values of comparison statements due to accidental inability of speakers to perceive or measure the compared entities). Significantly, the intuitive distinction between the two types of phenomena is not blurred, as they are shown to result from two different sources.

What is more, so far adherents of the epistemic approach to vagueness had to assume that the standards of vague adjectives are *inherently* unknowable. It is this assumption that many do not accept, doubting the reasons for this impossibility to know what the standards are. If they are out there in the world, why can't we discover them? Conversely, our proposal captures an intuitive distinction underling the difference between so-called 'inherent' and 'accidental' ignorance. If, as I claim, the ignorance underlying vagueness has a special *source*, then now we have an option to adopt an epistemic approach either without the assumption that vagueness is ignorance of a special *nature* ('inherent'), or with it, explaining the special *nature* based on the special *source* of truth value gaps.

Having presented these motivations, let us turn to my representation of the two different sources of truth value gaps.

3.2 Partial information about cutoff points vs. partial information about property values

Intuitively, a possible source of truth value gaps in the simple adjectival statements repeated below under (7) is vagueness/ ignorance concerning the interpretation of the predicates (e.g., *tall*), but the source of truth value gaps in the measure and comparison constructions repeated below under (8) is vagueness/ ignorance concerning the interpretation of the arguments (e.g., *Dan* and *Sam*). In this section I justify this claim and provide a representation for it.

(7) a. Dan is tall
b. Sam is intelligent
c. Mary is happy

(8) a. Dan is two meters tall
b. Dan is taller than Sam is

[3] In fact, while the above mentioned literature (mostly, Williamson's 1994 comprehensive assay on the epistemic approach) provides convincing arguments against a variety of objections to the epistemic approach, it seems to fail to provide convincing positive evidence (for a critical review of these arguments see Machina and Deuttsch, 2002). But very much the same can be said about the rival view (cf., e.g., Williamson's 2002 response to Machina and Deuttsch's arguments in favor of their own view).

c. Dan is twice as tall as Sam

Different completions t of an actual context c may resemble in the respects relevant for the evaluation of statements with vague predicates. For example, consider speakers who presume to be familiar with, e.g., Dan's height and Sam's IQ, but not necessarily with other property values, e.g., Dan and Sam's weight, hair color, mood, etc. The information carried by these speakers can be represented as a context c, whose completions t in T_c vary with regard to, e.g., the referent of the name *Dan* in them, but not with regard to the height of this referent (across T_c the individuals *Dan* denotes differ in weight, hair color, etc., but they all have the same height). Similarly, Sam's IQ (but not her weight or mood) may be context invariant, standing for the fact that, the given context c holds information about the IQ of *Sam*'s referent.

The statement *Dan is tall* is true in c iff this statement is true in every completion t of c, i.e. iff Dan's (completion-invariant) height exceeds the (completion-dependent) standard height (*tall*'s membership cutoff point) in all completions t of c; the statement is false in c iff Dan's height fails to exceed the standard height in each and every completion t of c, and the statement is undetermined otherwise (c.f., Kamp 1975). The standard truth conditions, then, can still do their old job (to my mind – good old job), regardless of the variance with regard to the individuals bearing the given height (Dan's different referents in T_c). Thus, even if Dan's height is completion-invariant in c, we can still say that a predicate like *tall* is vague (vaguer than a predicate like *bird*, *even number,* or *two meters tall*) in c, because the cutoff point of *tall* tends to vary across c's completions (or worlds), more than the cutoff point of *bird, even number* or *two meters tall* do.

While often the truth value gaps of statements like those in (7) pertain to the context dependency (variance across indices) of the cutoff point of adjectives, this cannot be the reason for truth value gaps in measure and comparison statements such as those in (8). Thus, we will now focus on these. Let us first consider the factors contributing to their truth conditions, and then consider our ability to gain information about these truth values.

Semanticists often associate gradable predicates like *long* with a mapping that is additive with respect to a dimension (Klein 1991). For example, the degree function of *long*, f_{long}, is *additive with respect to length*. Its values adequately represents the ordering, differences and ratios between quantities of length of individuals. Formally, the fact that the length of the concatenation (placing end to end) of any two individuals d_1 and d_2 (symbolized as $d_1 \oplus_{length} d_2$) equals the sum of lengths of the two separate individuals is represented by the fact that $f_{long}(d_1 \oplus_{length} d_2) = f_{long}(d_1) + f_{long}(d_2)$.

But there are infinitely many functions that are additive with respect to length, e.g., meter rulers and inch rulers represent two different additive mappings of entities to degrees (for instance, the numerical values assigned to the length of an A4 sheet of paper in inches and in centimeters are different). Furthermore, we cannot count directly quantities of the 'stuff' denoted by mass nouns such as *height*, *heat, happiness*, etc. These quantities have no known values (like 1,2,3,..) Thus, individuals d with a non-zero quantity of height (say, an A4 page) should be mapped to different numerals in different completions: $\exists t_1,t_2 \in T_c$: $f(tall,t_1)(d) \neq f(tall,t_2)(d)$.

Still, all rulers yield the same ordering and ratios between individuals' heights, and in any t, f(tall,t) represents these ordering and ratios. Given the additivity constraint, for every individual d, in every $t \in T_c$, individuals with n times d's height are mapped to the numeral $n \times f(tall,t)(d)$. The individual ordering is completion-invariant, then, and the ratios are completion-invariant numbers n. Thus, I propose that despite the fact that, e.g., f(tall,t) differs across completions in T_c, we have

information about the ordering and ratios between individuals' degrees of height, so there is no denotation-gap in predicates like *taller* (which relates to ordering), or *two meters tall* (which relates to ratios, meaning something like 'twice as tall as a meter'; cf. Sassoon 2009a,b). Based on this intuition, also in previous vagueness-based gradability theories (e.g., Kamp 1975; Fine 1975), the denotation of, e.g., *taller* does not vary across valuations in a vagueness-model.

So far we have shown that the denotation of measure and comparison predicates (e.g., *two meters tall* and *taller*) are fully specified (they are index-invariant). But, then, how come that sometimes we do not know the truth value of statements like *Dan is two inches tall* or *taller than Sam*? I submit that this is due to vagueness/ignornace pertaining to the interpretation of the arguments in these statements. Vagueness pertains to the interpretation of arguments even when they are proper names (e.g. *Dan* and *Sam*), or pronouns (as in *she is taller than two meters*).

In order to see that let us explicitly distinguish between the notion *discourse entities*, namely referents of proper names or terms in general, and *individuals*, namely indices in the domain of possible individuals D_x. Each discourse entity x corresponds to one and only one individual d in D_x in each world w or each completion t of any actual context c. As our information (knowledge) about each discourse entity is partial, it is consistent with it being any of indefinitely many individuals in D_x (cf., section 2 above).

The fact that I see a discourse entity x (say, Dan, that girl over there, you, the person approaching me, etc.), does not help me select the unique individual d in D_x that this discourse entity *is*. That is so because most of x's intrinsic properties are not accessible or hidden (x's precise height, size, weight, color, mood, etc.)

On the present proposal, when we do not know, e.g., the heights of two discourse entities, say, Dan and Sam, we do not know to which two individuals these names refer in c. When this happens, we may easily not know how their heights compare. If Dan's height is not accessible to me (its referent is 1.87m tall in t_1, 1.86m tall in t_2, etc.), I may not know whether *Dan is taller than Sam* it true or not, whether *Dan is 1.87 meters tall* is true or not, or whether *Dan is twice as tall as Sam* is true or not.

In fact, typically, formal semantic analyses build on the assumption that, e.g., facts pertaining to the height of all the entities in the domain are given in actual contexts c. If you measure them, you find out what they are, if you don't it does not have anything to do with the fact that measure and comparison statements have determined truth values. But in actuality, no speaker has access to, e.g., the height of all physical entities. In actual contexts c, the truth value of statements of the form *Dan is 85 inches tall* or *Dan is taller than Sam* may be unknown, if Dan and/or Sam are not accessible to the discourse participants whose knowledge c represents, or if their heights cannot be appropriately measured.

Recall that in 'conservative' representations of proper name interpretations (whereby proper names are rigid designators), individuals are often identified across different indices even if their intrinsic properties (height, shape, happiness extent, etc.) vary. On such theories, we may not know whether *Dan is taller than Sam* is true or not simply because we may be ignorant with regard to the intrinsic properties (e.g., heights) of the unique referents of names like *Dan* and *Sam*. This sounds very similar to the present proposal, but there is an important difference between the two accounts, as follows.

If we insist on this rigid-designation thesis (a rather common practice in linguistics, as I said earlier), we appear to be forced into the assumption that *tall*'s degree function may assign two entities d_1 and d_2 (e.g., the unique referents of names like *Dan* and *Sam*) values whose rank order

is different in different total contexts t in T_c, or whose ratio differs between total contexts. Only in virtue of that assumption would we be able to say that the interpretation of *Dan is taller than Sam* in c is undetermined. In one index $t_1 \in T_c$, the pair $<d_1,d_2>$ is an element of $[[taller]]_{t1}$, rendering the statement *Dan is taller than Sam* true in t_1), while in another index $t_2 \in T_c$, the reversed pair $<d_2,d_1>$ is an element of$[[taller]]_{t2}$, rendering the statement *Dan is taller than Sam* false in t_2.

However, in that case, we fail to account for the intuitive distinction between *tall* and *taller* (or *two meters tall*), i.e., between predicates that we intuitively regard as vague (admitting denotation gaps due to lack of sharp boundaries), and ones that we intuitively regard as sharp, i.e. ones that (due to their relatively sharp boundaries) rarely admit denotation gaps. Rigid designation for proper names means that in both predicate types the denotations would have to considerably vary across indices. In fact, as mentioned earlier, typically in vagueness-models (e.g., Kamp 1975), abstracting away from cases of accidental ignorance about property values, comparative and measure predicates are represented as possessing an index-invariant (non-vague) interpretation. The representation of truth value gaps in these cases, it is claimed, requires the incorporation of a representation of an epistemic modal base. But the question how a representation of epistemic modality may be incorporated into a vagueness model (or the other way around), remains unresolved, because if one actually combines the two representations, one risks losing the distinction between vague and non-vague predicates.

On the present proposal, this problem is solved. We account for the intuitive distinction between *tall* and *taller* (or *two meters tall*). We do so by representing the former as vague (admitting truth value gaps, due to lack of sharp boundaries), and the latter as non-vague (denoting the same set of individuals in every index t of T_c), e.g., if d_1 has more height than d_2 (in a given time point p), $<d_1,d_2> \in [[taller]]_t$ in every t in T (per p).

However, importantly, truth value gaps are nonetheless possible, in virtue of the index-dependent interpretations of the arguments of non-vague adjectives such as *taller* or *two meters tall*. While in each index predicates like *two meters tall* (or *taller*) denote the same set of individuals (or pairs of individuals), *they do not relate to the same discourse entities* (or pairs thereof), e.g., in indices in which the discourse entity Dan (the referent of *Dan*) is associated with d_1, and the discourse entity Sam (the referent of *Sam*) is associated with d_2, we regard the statement *Dan is taller than Sam* as true, and in indices in which it is the other ways around, we regard the statement as false, despite the fully determined interpretation of *taller* and due to the partial interpretation of *Dan* and/or *Sam*, i.e. our partial knowledge about their intrinsic properties (heights).

To wrap up, we have shown that besides phenomena that typically *can* be classified as instances of vagueness, we can also represent phenomena that are typically not classified as instances of vagueness, but of epistemic gaps (e.g., absence of knowledge concerning truth values of measure and comparison statements due to accidental inability of speakers to perceive or measure the compared entities). Significantly, the intuitive distinction between the two types is due to their two different sources. The phenomena typically not called ‘vagueness' in fact do not reflect absence of information (or knowledge) regarding the interpretation of a vague adjective like *tall* or a non-vague comparative like *taller*. Rather, they reflect absence of knowledge regarding the discourse entities (the interpretation of the arguments) in question.

We say *Dan,* but we do not know exactly to which possible individual we are referring, since we do not know all of Dan’s property values. *Dan’s* referent may be any one of several possible

individuals d in D_x that agree on many property values (e.g., their names, addresses, occupation, character, family, friends, and so on and so forth), but not on, e.g., their height. Truth value gaps in statements about measure and comparison, result from the fact that the denotation of a proper name may be empty, i.e. it may refer to different individuals differing with respect to their values in accidental intrinsic properties in different total contexts t in T_c of the context of evaluation c.[4]

For a complete formal model with graded individuals and partial interpretations for proper names as well as other terms see Sassoon (2007).

4. Conclusions

First, I adopt an assumption which is common in gradability theories (cf. Kennedy, 1999), namely that the ontology consists of a set of possible entities D, and a set of functions that map entities from D to real numbers depending on their values in gradable properties.

Second, I concur with Lewis (1986), taking different total contexts to be distinguished by the reference of proper names, but, unlike Lewis (1986), I take the domain of (possible) individuals to be common to all possible total contexts.

Third, in formal semantics, we usually take an individual to have a property iff it is in the denotation of that property. I say, in contrast, that an individual's 'real' properties are represented by property values (measures), i.e., individuals are distinguished by their property values (regardless where they fall with respect to the properties' cutoff points).

Fourth, I take different total contexts to also be distinguished by the cutoff points for vague-predicates in them.

This proposal captures the intuition that in different total contexts, the same individual, with the same property values, may still have different 'linguistic' or 'conceptual' (extrinsic, relational) properties, represented as belonging to the denotation of some vague predicate in one total context, but not in the other (it is only our interpretation of the vague word that has changed).

Fifth, individuals d in D are representations of possible real-world entities, represented by a completely given set of property values. Absence of information regarding a given property value (e.g., height) occurs only with regard to discourse entities, namely, referents of terms (when these referents are not yet specified), not with regard to individuals d in D.

This proposal captures the fact that even for predicates like *tall*, whose ordering relation depends on conventional well-known measuring systems, it is not the case that in actual contexts for any two referents of proper names we know whether they stand in the relation *taller* or not. At the same time, the interpretation of *taller* is fully specified. Individuals' heights are invariably and unambiguously determined, and, thus, we can always tell how their heights compare. However, when we use two proper names, we do not know exactly to which individuals in D they refer, and in particular, we may not know the heights of these discourse entities – so we may easily not know how their heights compare.

This proposal has the advantage that it allows regarding vagueness as merely ignorance, albeit ignorance about predicates' cutoff points, rather than about discourse entities' referents.

---

[4] The same can be said about other terms, such as *he* or *that* (cf. Sassoon, 2007, chapter 6). When we encounter and perceive an object (an actual worldly entity), only some of its properties are accessible to us, which means that we cannot tell which individual d in D it really is. This renders even variables or pronouns, whose interpretation appears to directly pick an individual (at least in deictic contexts), vague, in the sense that we do not know, and cannot know the exact d in D to which they refer.

References

Graff, Delia (2000) Shifting Sands: An Interest-Relative Theory of Vagueness, *Philosophical Topics* 28(1): 45–81.

Fein, Keith (1975), Truth, Vagueness and Logics, *Synthese* 30: 265-300.

Fox, Danny, and Martin Hackl. 2006. The universal density of measurement. Linguistics and Philosophy 29(5): 537–586.

Fox, Danny. 2007. Too many alternatives: density, symmetry and other predicaments. In Proceedings of SALT XVI, ed. M. Gibson and T. Friedman, Ithaca, NY: CLC Publications.

Kamp, Hans (1975) "Two theories about Adjectives", in Edward Keenan (ed.), *Formal Semantics for Natural Language*: Cambridge University Press, 123-155.

Kamp, Hans and Barbara Partee (1995), Prototype theory and compositionality. *Cognition* 57: 129-91.

Kennedy, Christopher (1999) *Projecting the adjective: The Syntax and Semantics of Gradability and Comparison*. NY: Garland. (1997 UCSC Doctoral dissertation). *Vagueness in Degree Constructions 487*

Klein, Ewan (1980), 'A Semantics for Positive and Comparative Adjectives', *Linguistics and Philosophy* 4: 1–45.

Klein, Ewan (1991) "Comparatives", in Arnim von Stechow and Dieter Wunderlich (eds.) *Semantik / Semantics, An International Handbook of Contemporary Research*, Berlin, New York: Walter de Gruyter, 673-691.

Krantz, David H.; Luce, R. Duncan; Suppes, Patrick, and Tversky, Amos (1971) *Foundations of measurement: Additive and Polynomial Representations*, San Diego, London: Uni. of Fribourg Academic Press.

Krifka, Manfred (2007) Negated antonyms: Creating and filling the gap. In Uli Sauerland and Penka Stateva (eds.), *Presupposition and implicature in compositional semantics*, Palgrave Studies in Pragmatics, Language and Cognition, Palgrave Macmillan: 163-177.

Landman, Fred. (1991). Structures for Semantics. Kluwer. Dordrecht.

Moltmann, Friederike (2006) "Comparatives without Degrees. A New Approach", A manuscript for the workshop on scalar meaning, University of Chicago.

Murphy, Gregory (2002), *The big book of concepts.* The MIT Press. Cambridge, MA.

Nouwen, Rick. (2008). Upper-bounded no more: the exhaustive interpretation of non-strict comparison, Natural Language Semantics 16: 271–295.

Russell, Bertrand, (1905) On Denoting. *Mind*, 14, 479-493. Repr. in Russell, Bertrand, *Essays in Analysis*, London: Allen and Unwin, 1973, 103-119.

Sassoon, W. Galit (2007) *Vagueness, Gradability, and Typicality – A comprehensive semantics analysis*, an unpublished doctoral dissertation, Tel Aviv University.

Sassoon, W. Galit (2009a) Negative predicates – The quantity metaphor and transformation values. Tova Friedman and Satoshi Ito (Eds.) *Proceedings of SALT XVIII*, CLC publications: 637-54 (peer reviewed; acceptance rate about 1/3).

Sassoon, W. Galit (2009b) Vagueness in degree constructions. In: Riester, Arndt; Solstad, Torgrim (eds.) *Proceedings of Sinn und Bedeutung 13*, Stuttgart.

Schwarzschild, Roger (2005) "Measure phrases as modifiers of adjectives", *Recherches Linguistiques de Vincennes* **35**, 207–228.

Stalnaker, Robert. (1978) ‘Assertion’. In Peter Cole (ed.) *Syntax and Semantics* 9: *Pragmatics* 315-332: Academic Press.
van Fraassen, Bas, C. (1969) "Presuppositions, Supervaluations and Free Logic", in K. Lambert (Ed.), *The Logical Way of Doing Things*: New Haven: Yale University Press, 67-91.
van Rooij, Robert (to appear) "Vagueness and Linguistics", in G. Ronzitti (ed.) *The Vagueness Handbook*, Dordrecht: Springer.
Veltman, Frank (1984) "Data Semantics", in JeroenGroenendijk, TheoJanssen and Martin Stokhof (Eds.), *Truth, Interpretation and Information Proceedings of the 3rd Amsterdam Colloquium*.
von Stechow, Arnim (1984a), "My Reaction to Cresswell’s, Hellan’s, Hoeksema’s and Seuren’s Comments", *Journal of Semantics* **3**, 183-199.
von Stechow, Arnim (1984b) "Comparing Semantic Theories of Comparison", *Journal of Semantics* **3**, 1-77.
Williamson, Timothy (1994) *Vagueness*, London and New York: Routledge.