Vagueness and Natural Language Semantics

Abstract

This chapter is devoted to the phenomenon of vagueness and the challenges that vague linguistic expressions raise for the kinds of semantic theories that are commonly used in descriptive and theoretical linguistics. The chapter aims firstly to show how we can study vagueness as an empirical phenomenon that can be observed in linguistic data; secondly, to outline why the observed properties of vague language are not easily accounted for in our classical semantic theories, and, finally, to describe a particular set of the responses to these challenges that are currently available in the formal semantics literature.

1 Introduction

This chapter is devoted to the phenomenon of vagueness and the challenges that vague linguistic expressions raise for the kinds of semantic theories that are commonly used in descriptive and theoretical linguistics. The puzzles and paradoxes raised by vague language (to be discussed below) have been extensively studied under many different angles in the fields of linguistics, philosophy, psychology and mathematics since antiquity. This chapter has three modest goals related to the project of developing a formal semantic theory for human languages: it aims firstly to show how we can study vagueness as an empirical phenomenon that can be observed in linguistic data; secondly, to outline why the observed properties of vague language are not easily accounted for in our classical semantic theories, and, finally, to describe a particular set of the responses to these challenges that are currently available in the formal semantics literature.

The chapter is laid out as follows: in section 2, we give a brief description of the properties of semantic theories commonly used in the formal semantics of natural languages when they are not concerned with vagueness specifically. Then, in section 3, we describe the empirical properties of vague predicates, focussing on three properties that cluster together with the class of relative and absolute adjectives in languages like English: borderline cases and borderline contradictions, fuzzy boundaries, and susceptibility to the Sorites paradox, and we outline how these properties challenge the class of semantic theories described in 2. With this in mind, in section 4, we outline some of the (many) available options for analyzing the puzzling properties of vague language, focussing on the frameworks that have received the most attention recently in natural language semantics. Section 5 concludes with some remarks about vagueness across categories and across languages. We provided some recommendations for further reading in section 6.

2 Our Classical Semantic Theories

Although, as will be discussed below, linguists have proposed many extensions in order to analyze the wide range of meanings that are assigned to the wide range of syntactic expressions found across languages, the formal systems commonly used in formal semantics almost uniformly have their basis in a Tarskian semantics for first order logic (FOL). It is therefore useful to take a moment to review this system, while highlighting the aspects that will be challenged by the existence of vague constituents.

2.1 Classical FOL

The language of FOL is defined as follows:

Definition 2.1. ***Vocabulary***. *The vocabulary of FOL consists of a series of individual constants* $a_1, a_2 \ldots$, *individual variables* $x_1, x_2 \ldots$, *unary predicate symbols* $P, Q, R \ldots$,[1] *quantifiers* $\forall$ *and* $\exists$, *and connectives* $\wedge, \vee, \neg$ *and* $\rightarrow$, *plus parentheses.*

Definition 2.2. ***Syntax***.

- *Variables and constants (and nothing else) are terms.*
- *If* t *is a term and* P *is a predicate symbol, then* $P(t)$ *is a well-formed formula (wff).*
- *If* ϕ *and* ψ *are wffs, then* $\neg\phi$, $\phi \wedge \psi$, $\phi \vee \psi$, $\phi \rightarrow \psi$, $\forall x \phi$, *and* $\exists x \phi$ *are wffs.*
- *Nothing else is a wff.*

Now we define the semantics for FOL. We first define models that consist of a set of individuals D and a function m.

Definition 2.3. ***Model***. *A model is a tuple* $M = \langle D, m \rangle$ *where* D *is a non-empty domain of individuals and* m *is a mapping on the non-logical vocabulary satisfying:*

- *For a constant* a_1, $m(a_1) \in D$.
- *For a predicate* P, $m(P) \subseteq D$.

The interpretation of variables is given by *assignments*.

Definition 2.4. ***Assignment***. *An assignment in a model* M *is a function* $g : \{x_n : n \in \mathbb{N}\} \rightarrow D$ *(from the set of variables to the domain* D*).*

A model together with an assignment is an interpretation.

Definition 2.5. ***Interpretation***. *An interpretation* $\mathcal{I}$ *is a pair* $\langle M, g \rangle$, *where* M *is a model and* g *is an assignment.*

We first associate an element from the domain D with every interpretation $\mathcal{I}$ and every term t.

[1] In this chapter, for simplicity, we will limit the discussion to systems with unary predicates because the $n-$ary predication case is simply a straightforward generalization of the unary predicate case.

Definition 2.6. ***Interpretation of terms.***

1. *If x_1 is a variable, then $\mathcal{I}(x_1) = g(x_1)$.*
2. *If a_1 is a constant, then $\mathcal{I}(a_1) = m(a_1)$.*

Finally, the satisfaction relation ($\vDash$) is defined as in definition 2.7.[2] In what follows, for an interpretation $\mathcal{I} = \langle M, g \rangle$, a variable x_1, and a_1 a constant, let $g[a_1/x_1]$ be the assignment in M which maps x_1 to a_1 and agrees with g on all variables that are distinct from x_1. Also, let $\mathcal{I}[a_1/x_1] = \langle M, g[a_1/x_1] \rangle$.

Definition 2.7. ***Satisfaction*** ($\vDash$). *For all interpretations $\mathcal{I} = \langle M, g \rangle$,*

1. $\mathcal{I} \vDash P(t)$ *iff* $\mathcal{I}(t) \in m(P)$
2. $\mathcal{I} \vDash \neg\phi$ *iff* $\mathcal{I} \nvDash \phi$
3. $\mathcal{I} \vDash \phi \wedge \psi$ *iff* $\mathcal{I} \vDash \phi$ *and* $\mathcal{I} \vDash \psi$
4. $\mathcal{I} \vDash \phi \vee \psi$ *iff* $\mathcal{I} \vDash \phi$ *or* $\mathcal{I} \vDash \psi$
5. $\mathcal{I} \vDash \phi \rightarrow \psi$ *iff if* $\mathcal{I} \vDash \phi$, *then* $\mathcal{I} \vDash \psi$
6. $\mathcal{I} \vDash \forall x_1 \phi$ *iff for every* a_1 *in* D, $\mathcal{I}[a_1/x_1] \vDash \phi$
7. $\mathcal{I} \vDash \exists x_1 \phi$ *iff there is some* a_1 *in* D, $\mathcal{I}[a_1/x_1] \vDash \phi$

In the next sections, we will discuss a number of theorems and arguments of FOL. Thus, we define the consequence relation between sets of formulas as follows:

Definition 2.8. ***Consequence*** ($\vDash$). *A set of formulas Ψ is a consequence of a set of formulas Φ (written $\Phi \vDash \Psi$) iff every interpretation which is a model of Φ is also a model of Ψ.*

- *Instead of $\{\psi\} \vDash \{\phi\}$, we will write $\psi \vDash \phi$.*
- *A formula ϕ is* ***valid*** *(written $\vDash \phi$) iff $\varnothing \vDash \phi$.*

More informally, given a multiple premise set Φ and a multiple conclusion set Ψ, $\Phi \vDash \Psi$ iff some member of Ψ is true when all members of Φ are true.

2.1.1 Aspects of FOL to Note

The first aspect of FOL that will become important in the discussion of vague language is that the two element Boolean algebra of truth values $\{0, 1\}$ (aka $\{true, false\}$) underlies definition 2.7 above. Well-formed formulas of FOL are mapped to exactly one of these these values in the way described by the definition of satisfaction. Crucially, **there are only two truth values** (this is known as the *Principle of Bivalence*. Additionally, **each interpretation of FOL is a (total) function:** it is both total and single-valued from the language into $\{0, 1\}$. Thus, the following (meta)theorems hold in FOL:

[2] As is common, we will use a_2 to refer to both the expression in the language and its interpretation ($\mathcal{I}(a_2)$), provided that it is clear from context which is meant.

(1) **No Gaps:**
For all $\mathcal{I}$, all $a \in D$ and all predicates P,
$\mathcal{I}(P(a)) = 1$ or $\mathcal{I}(P(a)) = 0$.

(2) **No Gluts:**
For all $\mathcal{I}$, all $a \in D$ and all predicates P,
It is not the case that $\mathcal{I}(P(a)) = 1$ and $\mathcal{I}(P(a)) = 0$.

Furthermore, definition 2.7 is recursive and truth-functional: which of the two truth values a wff is assigned is determined by the values assigned to its syntactic components. The components that are predicates are assigned a set of individuals. In the case of unary predicates, this structure is a set of individuals. **These sets have sharp boundaries**. For a given predicate denotation, an individual's degree of membership is either 0 or 1: in the set or out of the set. In this way, a unary predicate P naturally partitions the domain into the set of individuals included in P and its complement.

A final feature of FOL that is relevant for the puzzle of vagueness is the interpretation of negation. As shown in definition 2.7, a formula of the form $\neg P(a_1)$ is true just in case the corresponding formula $P(a_1)$ is false. In other words, $\neg P(a_1)$ is true just in case a_1 is in the complement of P in D. The partitioning nature of predicates and the definition of negation gives rise to certain validities in FOL (and related systems). In particular, there are no interpretations of FOL that can satisfy $\exists x_1(P(x_1) \wedge \neg P(x_1))$. This fact has an important effect on the semantic consequences that we can draw from such sentences. In particular, since no interpretations satisfy $\exists x_1(P(x_1) \wedge \neg P(x_1))$, by definition 2.8, any formula is a consequence of this sentence. In general, (3) follows immediately from the definitions given above.

(3) **Contradiction with Explosion:**
For all formulas ϕ, ψ,
$\{\phi, \neg\phi\} \vDash \psi$

Secondly, by virtue of the definition of negation, every individual must be in either the extension of a predicate P or its anti-extension $(D - m(P))$, and should not be in both. These are the laws of excluded middle (4) and non-contradiction (5).

(4) **Excluded Middle:**
For all predicates P,
$\vDash \forall x_1(P(x_1) \vee \neg P(x_1))$

In other words, all interpretations satisfy $P(a_1) \vee \neg P(a_1)$, for all $a_1 \in D$.

(5) **Non-contradiction:**
For all predicates P,
$\vDash \forall x_1 \neg(P(x_1) \wedge \neg P(x_1))$

In other words, all interpretations satisfy $\neg(P(a_1) \wedge \neg P(a_1))$, for all $a_1 \in D$.

Finally, we take a moment to highlight some other facts that hold in FOL given the semantics that we outlined above. Firstly, we can note that modus ponens is valid in FOL (6).

(6) **Modus Ponens:**
For all formulas ϕ, ψ,
$\{\phi \rightarrow \psi, \phi\} \vDash \psi$

Secondly we note that the deduction (meta)theorem holds (7).

(7) **Deduction Theorem:**
For all sets of formulas Γ, and all formulas ϕ, ψ,
If $\Gamma \cup \{\phi\} \vDash \psi$, then $\Gamma \vDash \phi \rightarrow \psi$.

Finally, we note that the consequence relation is transitive: (8) holds.

(8) **Transitivity:** :
For all formulas ϕ, ψ, χ,
If $\phi \vDash \psi$ and $\psi \vDash \chi$, then $\phi \vDash \chi$

As we will see in section 3, the semantic behaviour of vague predicates will appear to be in conflict with the picture described above.

2.2 Extensions in Linguistics

Clearly, the system just presented does not look very much like an interpreted grammar for English or any other possible natural language. And, we might wonder what bearing paradoxes for FOL (such as the Sorites paradox (to be described below)) might have on our theories of how meaning is constructed in human languages. However, within the Montagovian approach to the study of natural language semantics and pragmatics, the types of semantics that we give to grammars analyzing fragments of natural languages have much in common with the semantics of FOL described above, despite the many kinds of enrichments that linguists have proposed. For example, many advances in linguistic semantics have been made by proposing that the domain of individuals D is in fact sorted: it contains more than one kind of object, including (possibly) degrees (see section 4.4), events, worlds, times, numbers, among other things. The other type of domain enrichment common in linguistics is to impose additional relations between individuals that are not present in classical models for FOL (c.f. Link (1983), Keenan and Faltz (1985), Krifka (1989), and much later work). However, we can observe that the extensions proposed by linguists within the Montagovian tradition all preserve the properties that we highlighted above as being challenged by the phenomenon of vagueness.

1. Every interpretation function is still total, with $\{0, 1\}$ being the only truth values.
2. Constituents are still assigned sets. These sets may have more structure or consist of different sorts of objects than in many interpretations of FOL; however, the set-theoretic boundaries of these relations are still sharp.

3. In the vast majority of linguistic theories, negation is treated either as a propositional truth-reversing operator (ex. as in Chierchia and McConnell-Ginet (2000), Heim and Kratzer (1998) etc.) or as a more general complement operator (ex. Keenan and Faltz (1985), Winter (2001)). Thus, versions of excluded middle are taken to hold in natural languages as well.

4. Entailment in natural language is generally taken to have the same properties as in FOL (Heim and Kratzer (1998), Chierchia and McConnell-Ginet (2000) and every other textbook in formal semantics). Namely, semantic consequence is taken to be transitive, i.e. we want inferences like (9) to hold between natural language sentences, and some sort of deduction theorem should also hold (10).

(9) If *John came to the party early* $\vDash^{Eng}$ *John came to the party* and,
John came to the party $\vDash^{Eng}$ *John was at the party at some time*, then
John came to the party early $\vDash^{Eng}$ *John was at the party at some time.*

(10) *John came to the party early* $\vDash^{Eng}$ *John came to the party*, iff
$\vDash^{Eng}$ *John came to the party early* $\rightarrow$ *John came to the party*

In the next section, we will describe how the properties of predicates like *tall* appear difficult to accommodate in an analysis using the semantics of FOL; however, it should be clear based on the discussion above that these puzzles apply not only to simple classical first-order logical systems, but to the vast majority of semantic theories for natural language expressions that are proposed in linguistics and philosophy.

3 Diagnosing Vagueness

In this section, we present the three main characterizations of vague language in the sense relevant to semantics and discuss how the properties of vague language appear to be problematic for our classical semantic theories. These properties are the *borderline cases* property, the *fuzzy boundaries* property, and the *susceptibility to the Sorites paradox* property. We will first illustrate these properties and show how they cluster together with **relative** adjectives, such as *tall* and *friendly*, and then we will discuss the distribution of these properties with other kinds of adjectives.

3.1 Borderline Cases

The first characterization of vague predicates found in the literature, going back to Peirce (1902), if not earlier, is the *borderline cases* property. That is, vague predicates are those that admit borderline cases: objects of which it is unclear whether or not the predicate applies. Consider the following example with the predicate *tall*: If we take the set of American males as the appropriate comparison class for *tallness*, we can easily identify the ones that are clearly tall: for example, anyone over 6 feet. Similarly, it is clear that anyone under 5ft9" (the average) is not tall. But suppose that we look at John who is somewhere between 5ft9" and 6ft. Which one of the sentences in (11) is true?

(11) a. John is **tall**.
b. John is **not tall**.

For John, a borderline case of *tall*, it seems like the most appropriate answer is either "neither" or "both". In fact, many recent experimental studies on contradictions with borderline cases have found that the "both" and/or "neither" answers seem to be favoured by NL speakers (Ripley, 2011a; Alxatib and Pelletier, 2011; Serchuk et al., 2011; Egré et al., 2013). For example, Alxatib and Pelletier (2011) find that many participants are inclined to permit what seem like overt contradictions of the form in (12) with borderline cases, and Ripley (2011a) finds similar judgements for the predicate *near*.

(12) a. Mary is **neither tall nor not tall**.
b. Mary is **both tall and not tall**.

At first glance, we might hypothesize that what makes us doubt the principle of bivalence with borderline cases is that the context does not give us enough information to make an appropriate decision of which (of two truth values) the sentence *John is tall* has; for example, we are ignorant about John's height. However, as observed by Peirce, adding the required information does not make any difference to resolving the question: finding out that John is precisely 5ft11" does not seem to help us decide which sentence in (11) is true and which is false, or eliminate our desire to assent to contradictions for classical logical systems like (12).

Clearly, the existence of borderline cases poses a challenge for our classical semantic theories. As mentioned in the previous section, these systems are all bivalent: every sentence must have one of the two Boolean truth values. Thus, we have a puzzle.

3.2 Fuzzy Boundaries and Tolerance

A second characterization of vague predicates is the *fuzzy boundaries* property. This is the observation that there are (or appear to be) no sharp boundaries between cases of a vague predicate P and its negation. To take a concrete example: If we take a tall person and we start subtracting millimetres from their height it seems impossible to pinpoint the precise instance where subtracting a millimetre suddenly moves us from the height of a tall person to the height of a not tall person.

The fuzzy boundaries property is problematic for our classical semantic theories because we assign set-theoretic structures to predicates and their negations, and these sets have sharp boundaries. In principle, if we line all the individuals in the domain up according to height, we ought to be able to find an adjacent pair in the *tall*-series consisting of a tall person and a not tall person. However, it does not appear that this is possible.

Of course, one way to get around this problem would be to just stipulate where the boundary is, say, at another contextually given value for *tall*; however, if we were to do this, we would be left with the impression that the point at which we decided which of the borderline cases to include and which to exclude was arbitrary[3]. The inability to draw sharp, non-arbitrary

[3]This has been observed since Borel (1907/2014), see also Égré and Barberousse (2014).

boundaries is often taken to be the essence of vagueness (for example, by Fara (2000)), and it is intimately related to another characterization of vague language: vague predicates are those that are *tolerant*. Following's Wright (1975) (and his formulation), we will call a predicate **tolerant** with respect to a scale or a dimension Θ if there is some degree of change in respect of Θ insufficient ever to affect the justice Wright proposed this definition of vagueness as a way to give a more general explanation to the 'fuzzy boundaries' feature; however, versions of this idea have, more recently, been further developed and taken to be at the core of what it means to be a vague expression (ex. Eklund (2005), Smith (2008), van Rooij (2010), Cobreros et al. (2012)). This property is more nuanced than the 'fuzzy boundaries' property in that it makes reference to a dimension and to an incremental structure associated with this dimension, and it puts an additional constraint on what can be defined as a vague predicate: the distance between the points on the associated dimension must be sufficiently small such that changing from one point to an adjacent one does not affect whether we would apply the predicate. Immediately, we can see that *tall* is tolerant. There is an increment, say 1 mm, such that if someone is tall, then subtracting 1 mm does not suddenly make us call them not tall. Similarly, adding 1 mm to a person who is not tall will never make us call them tall. Since height is continuous, we will always be able to find some increment that will make *tall* tolerant. So, if we are considering very small things for whom 1 mm makes a significant difference in size, we can just pick 0.5 mm or whatever.

3.3 The Sorites Paradox

One of the reasons that vagueness has received so much attention in philosophy (in addition to linguistics) is that vague predicates seem to give rise to arguments, known as *sorites paradoxes*, that result in contradiction in FOL. Formally, the paradox can set up in a number of ways in FOL. A common one found in the literature is (13), where $\sim_P$ is a 'little by little' or 'indifference' relation[4].

(13) **The Sorites Paradox**

a. **Clear Case:** $P(a_1)$

b. **Clear Non-Case:** $\neg P(a_k)$

c. **Sorites Series:** $\exists a_1 \ldots a_n \forall i \in [1, n](a_i \sim_P a_{i+1})$

d. **Tolerance:** $\forall x \forall y((P(x) \wedge x \sim_P y) \rightarrow P(y))$

e. **Conclusion:** $P(a_k) \wedge \neg P(a_k)$

Thus, in FOL and other classical systems, as soon as we have a clear case of P, a clear non-case of P, and a Sorites series, through *universal instantiation* and repeated applications of *modus ponens* we can conclude that everything is P and that everything is not P. We can see that *tall* (for a North American male) gives rise to such an argument. We can find someone who measures 6ft to satisfy (13-a), and we can find someone who measures 5ft6" to satisfy (13-b). In the previous subsection, we concluded that *tall* is tolerant, so it satisfies (13-d), and, finally, we can easily construct a Sorites series based on height to fulfil (13-c). Therefore, we

[4]Note that, technically speaking, the Sorites argument is not stateable in the system that we set out above because the language does not contain binary predicates like $\sim_P$. Thus, the Sorites must be formulated in a slightly enriched language.

would expect to be able to conclude that this 5ft6" tall person (a non-borderline case) is both tall and not tall. We stress again that the Sorites is not only a paradox for Classical FOL. As discussed above, the semantic theories that linguists most commonly employ all assume bivalence and validate excluded middle, and modus ponens. Thus, the puzzles that vague predicates raise are widespread and shake the very core of the logical approach to natural language semantics.

3.4 Relative vs Absolute Adjectives

Although the vast majority of work done in semantics and philosophy of language has focussed on what are called **relative** adjectives like *tall*, we can observe similar (although not identical) properties with other classes of predicates. For example, what are called *absolute scalar* adjectives (predicates like *dry, wet, empty, straight, bent*, among others (Cruse, 1986; Kamp and Rossdeutscher, 1994; Yoon, 1996; Kennedy and McNally, 2005; Kennedy, 2007)) show a different pattern.

The first thing to observe about absolute predicates is that, as observed by (Pinkal, 1995; Kennedy, 2007, among others), these adjectives can sometimes be used precisely. For example, in contexts, when we use the predicate *straight*, we will want to pick out exactly those objects that are perfectly straight. A context that would favour the precise use of *straight* (which is discussed in Kennedy (2007)) is one in which we would say a sentence like (14).

(14) The rod for the antenna needs to be **straight**, but this one has a 1mm bend in the middle, so unfortunately it won't work.
(Kennedy, 2007, 25)

Most of the time, however, these predicates are not used in this way; rather, we can often use an absolute predicate to pick out individuals that deviate from the precise use of the predicate in some (contextually insignificant) way. For example, it is perfectly natural to say something like (15-a) even if there are some (insignificant) bends in the road. Likewise, depending on context, it may be natural to say something like (15-b) even if there are a couple of partiers at the club.

(15) a. This road is straight.
b. The nightclub is empty tonight.

Furthermore, we can observe that, in these 'loose' uses, predicates like *straight* or *empty* (what are known as *total* to *universal* absolute adjectives (Cruse, 1986; Kamp and Rossdeutscher, 1994; Yoon, 1996, and much subsequent work)) satisfy the tolerance principle. For example, suppose we want to go on a car trip, and one of us gets car sick very easily, so we only want to drive on straight roads. But, of course, it is not necessary for our purposes that the roads we drive on be perfectly straight; indeed this is most likely not possible. In this context then, we can pick '$\pm$ a 1mm bend' as an indifference relation for *straight*, because how could adding or subtracting a single millimetre bend make a difference to whether or not we would call a road *straight*? In this context then, (16) is true, and so *straight* appears to give rise to a sorites argument.

(16) For all roads x, y, if x is straight and x and y differ by a single millimetre bend, then y is straight.

Although the positive forms of total absolute adjectives can give rise to Sorites arguments, we can observe (following Égré and Bonnay (2010); Burnett (2012, 2014)) that these predicates are different from relative adjectives in that their negations (i.e. *not straight* or *not empty*) do **not** satisfy the tolerance principle. Suppose we pick **exactly** the same context (we want to go on a road trip; I don't want to get carsick. . .); therefore, ± a 1mm bend is still an indifference relations for *straight*. However, if we try to form a sorites argument with the negative form of sentences containing *straight*, we cannot. In particular, in the context described, (17) is false.

(17) For all roads x, y, if x is not straight and x and y differ by a single 1mm bend, then y is not straight.

In particular, the appropriate counter-example is the case when we move from a road with a 1mm bend (i.e. a road that is *not straight*) to a perfectly straight road (i.e. a road that is not *not straight*). More generally, unlike relative adjectives, absolute adjectives display certain non-symmetries in their judgments of indifference; that is, although, depending on the context, we might consider an object that is not perfectly straight to be straight (and we do this all the time), we will never consider an object that is perfectly straight to be not straight.

Thus, we see a first difference in Sorites susceptibility between the relative and total absolute adjectives. We can see another such difference when we compare total predicates with another subclass of absolute adjectives: what are called *partial* (or *existential*) absolute scalar adjectives (ex. *wet, bent, sick, dirty* etc.). Unlike total predicates, these adjectives give rise to a Soritical argument when used in negative sentences. For example, suppose I am getting out of the shower and looking for a towel to dry myself with. I need to pick a towel that is not wet; however, it doesn't really matter if there are a couple of drops of water on it. Thus, in this context, we can pick the relation ± one drop of water as an indifference relation for *wet*, and *wet* satisfies the negative version of the tolerance principle (18).

(18) For all towels x, y, if x is not wet, and x and y differ by a single drop of water, then y is not wet.

This time, however, it is the positive form of the predicate that is not tolerant: (19) is false, and I again highlight the existence of non-symmetry in judgments of indifference with these predicates: although, depending on context, we might consider an object that has one drop of water on it to be not wet (and we frequently do), we will never consider a bone-dry object to be wet.

(19) For all towels x, y, if x is wet and x and y differ by a single drop of water, then y is wet.

In summary, we see a diverse set of fine-grained patterns of sorites-susceptibility within the adjectival domain in languages like English. In the rest of the paper, however, we will limit our attention to the previously proposed solutions to the challenges posed specifically by relative adjectives; however, see (Pinkal, 1995; Kennedy, 2007; Toledo and Sassoon, 2011; Burnett, 2014, among others) for extensions of contextualist, epistemicist and multi-valued accounts of vagueness to absolute predicates.

3.5 Higher-order vagueness

A large topic that we will be unable to do true justice to in this chapter is higher-order vagueness (HOV). Part of what makes higher-order vagueness complex, is that there are arguably different phenomena that could be characterised as evoking vagueness of a higher-order. In no particular order, some of these are detailed below. These are clearly not all independent from each other, but tend to lead to the framing of related questions with a different emphasis.[5] In Section 4, we will occasionally highlight how different semantic accounts of vagueness fair with respect to these different conceptions of HOV.

Lexical HOV. If a semantic theory captures vagueness of first order natural language predicates (e.g. *tall* and *green*), does the same theory get the right results for second-order (or n^{th}-order) *natural language* predicates such as *really* and *very* (applied to relative adjectives), *completely* and *totally* (applied to absolute adjectives), *definitely*, *truly* and *certainly* (as VP modifiers). As long as, for example, combining the semantics of first and higher order predicates makes the right predictions, even for iterations of applications (*really tall*, *definitely (is) really tall* etc.), addressing Lexical HOV can be a relatively bounded enterprise.

Formal HOV: One hallmark of vagueness is that there is a tension in identifying the boundary between a vague predicates positive extension and its negative extension. If a semantic theory accommodates vagueness by discriminating one or more areas in the extension of a predicate beside those for which the predicate is true *simpliciter* or false *simpliciter*, what can be said about the boundary between the positive (or negative) extension and the intermediate areas. For example, if a sharp cut-off between positive extension $\mathcal{I}(P^+)$ and $\mathcal{I}(P^-)$ is assuaged by the introduction of an intermediate extension $\mathcal{I}(P^\pm)$, should we feel concerned if there are sharp cut-off points between $\mathcal{I}(P^+)$ and $\mathcal{I}(P^\pm)$ and between $\mathcal{I}(P^\pm)$ and $\mathcal{I}(P^-)$?

Formal HOV is usually framed in terms of a formal Δ (definiteness) operator. For example, if $\mathcal{I}(P) = \{a_1, a_2, a_3\}$ and $\mathcal{I}(\neg P) = \{a_4, a_5, a_6\}$, then the sharp cut-off between P and $\neg P$ can perhaps be assuaged by thinking about what is definitely P and definitely not-P. For example, if $\Delta(P)$ is $\{a_1, a_2\}$ and $\Delta(\neg P)$ is $\{a_5, a_6\}$, then at least there isn't a sharp cut-off between what is definitely P and definitely not-P. However, this yields no respite, since there is now a sharp cut-off between $\Delta(P) = \{a_1, a_2\}$ and $\neg\Delta(P) = \{a_3, a_4, a_5, a_6\}$. Asking what is definitely definitely P ($\Delta(\Delta(P))$) may remove the sharp cut-off at the second order, but introduces one at the third (between $\Delta(\Delta(P))$ and $\neg\Delta(\Delta(P))$), and so on.

Metasemantic HOV: What we call here *metasemantic* HOV frames issues of HOV in terms of e.g, justification, entitlement, and correctness of e.g., belief or assertion. The problem of HOV, in this form, is discussed at length by Wright (1975). In simple terms, the problem is the following: No matter what a semantic theory says about the extension of P (or the truth conditions of P), are we still left with a problem of under what conditions one would be

[5] It should however be noted that some have questioned whether or not HOV even exists (Wright, 2009) or whether first-order vagueness is, in fact, higher-order vagueness (Bobzien, 2015).

justified/correct/entitled to assert/believe P? For example, if a theory of vagueness results in a completely smooth transition between P and $\neg P$ (be it in terms of degrees of truth or something else), are we left with any answer to which point one should cease to use P in, for example a P-based sorites series? In other words, even if we get vague truth conditions right, does a related problem arise for e.g., correct or justified use-conditions? In other words, vagueness is arguably about blurriness and/or borderline cases, but agents must sometimes apply a predicate, not apply a predicate (or hedge). Two options seem to be available: (A) Develop a semantics which has sharp boundaries. This allows an easy mapping between when it is right to apply P (or not, or hedge), but does not necessarily make for a satisfactory theory of vagueness. (B) Develop a semantic theory of vagueness that provides graded denotations of predicates. This arguably gives a better account of the truth conditions of vague predicate, however, it cannot straightforwardly map semantics into, say, a theory of correct assertion. Furthermore, if we pick option (B), and then try to give a theory of e.g., correct/justified assertion, then we seem to be forced to either (B1) pick a point on the graded scale as the last point at which one can correctly assert P, or (B2) endorse a graded view of e.g. justified assertion. If we choose (B1), we arguably have an unsatisfactory account (a small difference in a object can make a big difference in whether we are justified in applying P). However, if we pick (B2), we have just shifted the problem one level up (to when we are e.g. *truly* justified in applying P).

On this conception of HOV, giving an account of vagueness looks like a matter of deciding at which level it is acceptable to have sharp cut-off points, since massaging them away at one level seems to force them to reappear at another.[6]

4 Major Approaches to the Analysis of Vague Predicates

In section 2, we set out an overview of a simple semantics based on CFOL which included a number of principles and theorems of this system. A good way to understand the plethora of theories of vagueness is to see exactly where each theory departs from this classical FOL foundation. In this section, we have selected some exemplars of each approach, and detail what divergences from the classical position each makes, the consequences of doing so, and a few outstanding challenges each type of account faces.

We will proceed in a semi-chronological order so that we may detail how the challenges with earlier approaches to analysing vagueness led to further developments. We will also group approaches together, as far as possible in terms of their semantic similarity, however, this will occasionally interfere with the chronological ordering.

Although there are still defenders of versions of all of the approaches to be detailed, there is a pattern with respect to the dominance, or at least, prominence, that some theories have had over the decades. In the 1960s and 70s, Fuzzy logical approaches were developed (§4.1. These depart significantly from classical approaches. The problems with such radical departures from classicism lead, in the 70s to non-classical theories that retain classical theorems such as supervaluationism and some forms of contextualism (see §§4.2–4.3)). In the 90s and 2000s, attempts were made to remain entirely classical in the form of epistemicism, degree-

[6]See Sutton (2017) for further discussion with a focus on probabilistic approaches to vagueness.

based semantics, and probabilistic accounts (§4.4). More recently, however, further departures from classicism have been suggested such as the *Tolerant, Classical Strict* approach (TCS) (see §4.5), and defences of new versions of older positions have been made (such as more sophisticated fuzzy logical approaches).

4.1 Fuzzy Logical Approaches

Formal systems for multiple valued logics (including infinite-valued logics) were first developed by Łukasiewicz (1922/1970). However, following later (independently developed) work into fuzzy sets (Zadeh, 1965) and the development of a logic based on fuzzy sets (Goguen, 1969), proposals for analysing vague predicates have been made within a fuzzy-logical framework. The main motivation for fuzzy-logical approaches is twofold. Vagueness seems fundamentally to be (i) a matter of degree, and (ii) a failure of bivalence. A natural thought, therefore, is to replace standard classical set theory with a set theory based on degree where an element is only a member of a fuzzy set to some degree in the range $[0, 1]$. In terms of logic, this replaces the two classical truth values $\{0, 1\}$ with a range of degrees of truth $[0, 1]$, and connectives are degree-of-truth-functional. There is room for alternatives in defining connectives, but those standardly presented are for negation, conjunction, and disjunction. There have also been different suggestions for fuzzy conditionals. The original suggestion from Łukasiewicz (1922/1970) is given below (he was also the first to suggest the negation rule).[7]

Definition 4.1. ***Interpretation of Fuzzy Logical Propositions.***
For an interpretation function for fuzzy values $\mathcal{I}_f$, and for all propositions ϕ, ψ:

1. $\mathcal{I}_f(\phi) \in [0, 1]$
2. $\mathcal{I}_f(\neg\phi) = 1 - \mathcal{I}_f(\phi)$
3. $\mathcal{I}_f(\phi \wedge \psi) = min\{\mathcal{I}_f(\phi), \mathcal{I}_f(\psi)\}$
4. $\mathcal{I}_f(\phi \vee \psi) = max\{\mathcal{I}_f(\phi), \mathcal{I}_f(\psi)\}$
5. $\mathcal{I}_f(\phi \to \psi) = \begin{cases} 1, & \text{if } \mathcal{I}_f(\phi) \leq \mathcal{I}_f(\psi) \\ 1 - \mathcal{I}_f(\phi) + \mathcal{I}_f(\psi) & \text{if } \mathcal{I}_f(\phi) > \mathcal{I}_f(\psi) \end{cases}$

The degree of truth of a proposition is inversely proportional to the degree of truth of its negation. A conjunction can only be as true as its least true conjunct. And, a disjunction is as true as its truest disjunct. Such a fuzzy system is classical if restricted to values at the limit, but an inclusion of the continuum of values makes fuzzy systems depart from classical logic in many other ways. The Principle of Bivalence clearly fails since the truth values of propositions are a continuum. It should be evident that neither excluded middle (4) nor non-contradiction (5) are preserved in a fuzzy system. Assuming that $\vDash$ in a fuzzy system preserves absolute truth, propositions of the forms $\phi \vee \neg\phi$ and $\neg(\phi \wedge \neg\phi)$ cannot be theorems. Values for such propositions fall in the range $[0.5, 1]$, and so represent up to a 0.5 drop in truth value.

With respect to the sorites series, there is room for interpretation within a fuzzy system, specifically on how the conditional is defined, and whether validity is defined in terms of

[7] For a useful overview of fuzzy connectives, please see Smith (2008, §2.2.1).

preservation of absolute truth or preservation of degree of truth. Common to all approaches, however, is that the clear case premise is perfectly or near perfectly true, the clear non-case premise is perfectly, or near perfectly false, and values for intermediate cases form a gradation in between. If the sorites is viewed as a series of applications of *modus ponens*, then the conclusion of each step is marginally falser than at the previous step.

This being said, fuzzy logics have been widely criticised (Kamp, 1975; Williamson, 1994; Edgington, 1997, among many others). Most objections to fuzzy approaches derive from one feature of fuzzy systems that is hard to swallow. Flat contradictions receive values in the range $[0, 0.5]$. Many have reacted negatively to the idea that a flat contradiction can be anything other than completely false. Yet, things are worse for fuzzy logic than that. As we saw in section 3.1, there have been empirical observations that "F and not F" responses to borderline cases of vague predicates are common. Perhaps, then, flat contradictions needn't always be false as the philosophical orthodoxy would suggest. However, taking these diverging intuitions at face value, either a flat contradiction should be valued as totally false, or, in borderline cases, as totally acceptable/true. Unfortunately for fuzzy logic, flat contradictions made about central borderline cases receive values of 0.5 which satisfies neither intuition. That said, one could defend a view that intuitions regarding flat contradictions track *acceptability* as opposed to limit-value degrees of truth. If this were so, then, provided that, for example, the acceptability of an outright assertion of a flat contradiction could be, say 0 when its degree of truth value is 0.5, then some intuitions can be accommodated. (See e.g. Smith (2008) for such a proposal.) Ultimately, the success of such a position will turn on defending the view that intuitions surrounding contradictions track acceptability and not (absolute) falsity. Other challenges arise, too. For example, a sorites argument could now be formulated surrounding acceptability, namely a metasemantic HOV problem (see Wright (1975) for discussion, Smith (2008) for a reply, and Sutton (2017) for discussion of difficulties with proposals such as Smith's).

However, a recent development of a fuzzy approach (Alxatib et al., 2013) has embraced the idea that flat contradictions can be completely true. This is achieved by a scaling operation. Simplifying somewhat, Alxatib et al. define intensional connectives. For example, the value of an intensional conjunction $\owedge$ is calculated in terms of $\mathcal{I}_f(\phi \wedge \psi)$, but also the *floor value* (the minimal possible value for the fuzzy conjunction) $\mathbf{f}(\phi \wedge \psi)$, the *ceiling value* (the maximal ceiling possible value for the fuzzy conjunction) $\mathbf{c}(\phi \wedge \psi)$:

(20)
$$\mathcal{I}_f(\phi \owedge \psi) = \begin{cases} \mathcal{I}_f(\phi \wedge \psi) & \text{if } \mathbf{f}(\phi \wedge \psi) = \mathbf{c}(\phi \wedge \psi) \\ \dfrac{\mathcal{I}_f(\phi \wedge \psi) - \mathbf{f}(\phi \wedge \psi)}{\mathbf{c}(\phi \wedge \psi) - \mathbf{f}(\phi \wedge \psi)} & \text{otherwise} \end{cases}$$

The floor and ceiling of $\mathcal{I}_f(\phi \wedge \neg\phi)$ are 0 and 0.5 respectively, so for an absolute borderline case the proposition $\phi \owedge \neg\phi$ receives a value of 1 ($0.5 - 0$ divided by $0.5 - 0$).

However, a further species of problem has been raised against fuzzy logic based accounts. The following is based on an example in Edgington (1997). Suppose that, as a contingent matter, $\mathcal{I}_f(F(a)) = 0.5$ and $\mathcal{I}_f(F(b)) = 0.4$. This means that a is a little more F than b. If a is a little more F than b, then it should not be the case that b is F when a is not, thus the proposition $F(b) \wedge \neg F(a)$ must have a perfectly false (or at least very low truth value). However, on a fuzzy system, such a proposition is far from perfectly false: $\mathcal{I}_f(F(b) \wedge \neg F(a)) = min\{0.4, 1-0.5\} = 0.4$.[8]

Furthermore, this species of counterargument arguably applies to intensional fuzzy conjunction. The conjunction of two different propositions should on Alxatib et al.'s independent logic system have a floor of 0 and a ceiling of 1 (since the F-ness of a and b is contingent), in which case $\mathcal{I}_f(\phi \otimes \psi)$ reduces to $\mathcal{I}_f(\phi \wedge \psi)$. We still do not get a low enough value.

The difficulty of capturing logical dependencies between propositions ("truths on a penumbra" in Fine (1975)), led some to move further towards classicism. Supervaluationism, for instance, retains fuzzy logic's rejection of bivalence, but does not adopt degree functionality.[9]

4.2 S'valuationism

S'valuationism is a coverall term for two related but distinct approaches: supervaluationism and subvaluationism.[10] Both are associated with a semantic analysis of vagueness in that vagueness is characterisable in terms of the meanings of terms. The principle adjustment S'valuationism makes to the classical approach is to drop the total interpretation function assumption. Vague predicates are vague because their interpretations underdetermine their extensions. However, a problem with truth value gaps in partial models is that one must decide how to interpret logical constants. If a is in the extension gap of P, what value should be given to, for example $P(a) \wedge \neg P(a)$ and $P(a) \vee \neg P(a)$?

Both s'valuationisms adopt a similar strategy in answering this question, but differ on the final valuation rule. Partial models determine for some elements of the domain whether they are or are not in the extension of a predicate P. Call those that are, *the positive extension* (${}^{+}P$), and those that aren't, *the negative extension* (${}^{-}P$). The basic idea is then that partial models can be extended to be classical models (interpretation functions can be made total). The extension of a partially interpreted predicate to a classically interpreted one is called a *precisification*. For a partial model, there will usually be more than one way to extend/precisify it. For example, if $a, b \in {}^{+}P$, $e, f \in {}^{-}P$, and c, d are in the extension gap of P, then there are multiple ways to extend the model ($\mathcal{I}_{s_i}$ is a total extension of the partial interpretation function $\mathcal{I}_s$):

(21)
$$\mathcal{I}_{s_0}(P) = \{a, b\}$$
$$\mathcal{I}_{s_1}(P) = \{a, b, c\}$$
$$\mathcal{I}_{s_2}(P) = \{a, b, c, d\}$$

Since functions $\mathcal{I}_{s_i}$ are total, they sort all members of the domain into either the extension of the anti-extension of a predicate. S'valuationsim restricts possible extensions of models in two ways. First, no extension can shift an object from the positive extension of a predicate into it's

[8]A parallel problem arises for material implication since $\mathcal{I}_f(F(b) \supset \neg F(a)) = \mathcal{I}_f(F(b) \supset F(a)) = 0.6$. On a non-material definition of the conditional, such as Łukasiewicz's conditional (5 in Definition 4.1), we arguably get a worse result, since $\mathcal{I}_f(F(b) \rightarrow \neg F(a)) = \mathcal{I}_f(F(b) \rightarrow F(a)) = 1$.

[9]In Section 4.4 we shall consider accounts that incorporate degrees into semantics and maintain bivalence. Also, it is possible to incorporate degrees into a supervaluationist model (Williams, 2011).

[10]The term is first used by Ripley (2013b) and Cobreros et al. (2011).

classically evaluated anti-extension. For example, both of the following would be inadmissible interpretations:

(22)
$$\mathcal{I}_{s_{\#1}}(P) = \{a\}$$
$$\mathcal{I}_{s_{\#2}}(P) = \{a, b, c, d, e\}$$

The second restriction turns on what Fine (1975) calls *penumbral connections*. Assume that $\langle a, b, c, d, e\rangle$ are individuals ordered in terms of decreasing P-ness. One cannot include in $\mathcal{I}_{s_i}(P)$ an individual that is less P than an individual not included in $\mathcal{I}_{s_i}(P)$. For example, the following would be an inadmissible precisification:

(23)
$$\mathcal{I}_{s_{\#3}}(P) = \{a, b, d\}$$

S'valuationisms then compute the S'value of a proposition in terms of all (supervaluationism) or some (subvaluationism) of these classical extensions to partial models.

4.2.1 Supervaluationism

Pioneered as an approach to vagueness by Mehlberg (1958) but brought to more prominent attention by Fine (1975), Kamp (1975) and Kamp and Partee (1995), a proposition is true iff it is true on all valuations and false iff it is false on all valuations.[11] It is neither true nor false if it is not true and not false. Supervaluationism is an advance on merely "gappy" accounts since one can supervaluate complex propositions as well as atomic ones.

All classical theorems are valid on a supervaluational approach (however the consequence relation differs in a way to be specified shortly). For example, instances of excluded middle are all true because they are true on all classical precisifications ($\vDash_{\text{superval}} \phi \vee \neg\phi$). Instances of contradiction are all false because they are false on all classical precisifications ($\vDash_{\text{superval}} \neg(\phi \wedge \neg\phi)$).

One key difference with classical models is that the metasemantic Principe of Bivalence does not hold in a supervaluationist system, since a proposition can, when supervaluated, be neither true nor false (i.e. true on some admissible precisifications and false on others). That is to say that, as a gappy theory, supervaluationism is weakly paracomplete (Hyde, 2008). Classical logic doe not distinguish between the following consequences. For any proposition ψ, one can conclude the multiple conclusion $\phi, \neg\phi$ (at least one of ϕ and $\neg\phi$ is entailed), or that $\phi \vee \neg\phi$ is true.

(24)
$$\psi \vDash_{\text{CL}} \phi \vee \neg\phi$$
$$\psi \vDash_{\text{CL}} \phi, \neg\phi$$

One way to think about this is that semantic models based on CL have only one valuation (a classical one). Hence, for any valuation in which $\phi \vee \neg\phi$ holds, the same valuation will mean that either ϕ or $\neg\phi$ holds (see 4 in Definition 2.7). In contrast, supervaluationist consequence supports one but not the other:

(25)
$$\psi \vDash_{\text{superval}} \phi \vee \neg\phi$$
$$\psi \nvDash_{\text{superval}} \phi, \neg\phi$$

[11] There is also an early outline of a degree based notion of supervaluationism in Lewis (1970).

The reason for this is that on every classical valuation, it must be true true that $\phi \vee \neg\phi$, however, since we are now working with truth across all classical valuations, it is possible that neither ϕ nor $\neg\phi$ are true across all valuations, hence one cannot conclude that one of ϕ and $\neg\phi$ are true.

Edgington (1997, p. 310) provides some natural language examples of when such inferences to go through, albeit in defence of her *verities* view (§4.4.3). For example, "A library book can be such that it is not clear whether it should be classified as Philosophy of Language or Philosophy of Logic; but if we have a joint category for books of either kind, it clearly belongs there." There is, however, some debate about whether such examples are persuasive (Hyde, 2008, ch. 4).

With respect to sorites arguments, supervaluationism can answer why there is no sudden transition from true instances of a predicate to false ones, because there are many cases in between of which the predicate is neither true nor false. Supervaluationism deems sorites arguments valid but unsound. The false premise is the tolerance premise. So for any variable assignment:

(26) $$\mathcal{I}_{\text{superval}}(\forall x \forall y((P(x) \wedge x \sim_P y) \to P(y))) = 0$$

On every classical precisification, there is a false instance to the premise (making it false), so therefore the tolerance premise is supervaluated as false. At least some of the instances of the tolerance premise (the tolerance conditionals) are neither true nor false (they are true/false on some but not all classical precisifications).

This diagnosis of the sorites has a down side, however. The most prominent objection to supervaluationism is that, although the falsity of the tolerance premise might seem appealing, its negation is true on all valuations ($\mathcal{I}_{\text{superval}}(\exists x \exists y((P(x) \wedge x \sim_P y) \wedge \neg P(y))) = 1$). Yet this existential premise can be interpreted as saying, counterintuitively, that vague predicates have sharp boundaries. See Hyde (2008, ch. 4) for a review of the different supervaluationist reactions to this problem.

4.2.2 Subvaluationism

Subvaluationism (defended, for example, in Hyde (1997), and more recently in Hyde and Colyvan (2008); Cobreros (2011)) has, perhaps until lately, received a good deal less attention that its supervaluational sister. On subvaluationism, a proposition is true if it is true on at least one admissible classical precisification, and false if it is false on at least one admissible classical precisification (alternatively, a proposition is true if not all classical interpretations make it false). In other words, subvaluationism is the dual of supervaluationism.

Unlike supervaluationism, subvaluationism has no truth value gaps ((1) holds). Since every proposition is either true or false on every precisification, every proposition will be either true or false on at least one precsification. Instead, however, we get truth value *gluts* ((2) does not hold). Whereas the supervaluational truth function is partial (some statements are supervaluated as neither true nor false), the subvaluational truth 'function' is not properly speaking a function at all, since it assigns more than one value to some propositions.

All classical theorems that are supervaluationally valid are subvaluationally valid, albeit for different reasons than with supervaluationism. For example, instances of excluded middle are all true because they are not *false* on any classical precisifications. Instances of non-contradiction are all false because they are not *true* on any classical precisifications. Although Non-Contradiction holds ($\vDash_{\text{subval}} \neg(\phi \wedge \neg\phi)$), the semantic equivalent of non-contradiction fails, namely, it is not true on subvaluationism that no proposition is true and false. However, the extent to which subvaluationism is paraconsistant is constrained. Subvaluationism is *weakly paraconsistant*. For classical logic, both a single premise contradiction and a set of inconsistent premises lead to explosion ((3) holds):

(27)
$$\begin{aligned} \phi \wedge \neg\phi &\vDash_{\text{CL}} \varnothing \\ \phi, \neg\phi &\vDash_{\text{CL}} \varnothing \end{aligned}$$

However, subvaluationism distinguishes the assertion of a contradiction ($\phi \wedge \neg\phi$) from a classically inconsistent set of premises (e.g., $\{\phi, \neg\phi\}$):

(28)
$$\begin{aligned} \phi \wedge \neg\phi &\vDash_{\text{subval}} \varnothing \\ \phi, \neg\phi &\nvDash_{\text{subval}} \varnothing \end{aligned}$$

In other words, (3) does not hold. This follows because on all classical precisifications, every statement of the form $\phi \wedge \neg\phi$ is false, hence for no statement of the form $\phi \wedge \neg\phi$ is it the case that some classical precisifications are true and others false. However, for some statement ϕ, it may be the case that ϕ is true on some precisifications, but false on others (therefore ϕ can be both subvaluationally true and subvaluationally false).

Given this feature of subvaluationist logic, the classical and subvaluationist consequence relations diverge with respect to conjunction introduction:

(29)
$$\begin{aligned} \phi, \psi &\vDash_{\text{CL}} \phi \wedge \psi \\ \phi, \psi &\nvDash_{\text{subval}} \phi \wedge \psi \end{aligned}$$

One thing that comes as an immediate benefit of adopting a subvaluationist logic is that it captures some of the empirical data that supervaluationism cannot, namely, that it seems very natural, for borderline cases of applying vague predicates, to say that something is both P and not-P. Nonetheless, there is an anomaly. If a is a borderline case of F, then a subvaluationist can say that 'a is F' is true and 'a is F' is false. However, it seems just as natural to express this as a is F and not F. But normally, this would be modelled as the proposition $F(a) \wedge \neg F(a)$ which is subvaluationally false! So the subvaluationist has to engage in something akin to doublethink when describing borderline cases, or, at least deny that "a is F and not F" expresses the proposition $F(a) \wedge \neg F(a)$.[12]

[12] We do not wish to be unduly unfair on subvaluationism. A similar anomaly actually infects a number of approaches to vagueness, except when applied to Bivalence and excluded middle. For example, if a a borderline case of F, supervaluationists can assert $F(a) \vee \neg F(a)$, but cannot assert that Either $F(a)$ is true or $F(a)$ is false. Similarly, on a probabilistic *verities* based approach (see section 4.4.3), it is perfectly true that $F(a) \vee \neg F(a)$, but it may be far from clearly true that $F(a)$ and far from clearly true that $\neg F(a)$.

The major stumbling block that subvaluation faces, however, is perhaps more cultural and historical. The dominant philosophical influence in semantics is the Russell-Carnap-Quine tradition which *de facto*, even if not *de jure*, makes a dialethic position such as Subvaluationism harder to convince people of. An exemplar of this conservative stance towards the impact of vagueness on logic is Williamson (1994) (but also see Sorensen (1988, 2001)). For more in-depth discussion of defences of subvaluationism, see Hyde (2008, ch. 4).

4.2.3 Higher-Order Vagueness in S'valuationism

Both forms of S'valuationism face the challenge of higher-order vagueness. Predicates such as "tall" are vague, but predicates such as "clearly tall"/"truly tall" are also vague. Yet the partial models which are extended in this framework are most intuitively interpreted as determining the clear extensions of predicates. However, if this is so, then clearly/truly-P comes out as non-vague, since it will be interpreted either as the set of entities in P under every precisification (supervaluationism), or as the set of entities not in $\neg P$ under any precisification (subvaluationism). However, under other renderings of higher-order vagueness problems (such as those based on "gap principles"), subvaluationism has been argued to outperform both supervaluationism and classicism (Cobreros, 2011).

4.3 Contextualism

4.3.1 An exemplar of a contextualist account: (Kamp, 1981)

An early contextualist approach to vagueness is explored in Kamp (1981). Kamp, who originally defended supervaluationism (Kamp, 1975), became dissatisfied with a consequence of the supervaluationist treatment of the sorites mentioned in 4.2.1. Namely, the falsity of Tolerance (the truth of $\neg\forall x\forall y((P(x) \wedge x \sim_P y) \rightarrow P(y)))$ implies $\exists x\exists y((P(x) \wedge x \sim_P y) \wedge \neg P(y))$ which could reasonably be interpreted as a denial of the vagueness of P.

Kamp's novel suggestion was to add a restriction on the truth of a universal sentence that means that it can be false without the equivalent existential sentence being true.[13] This is achieved via a complex account of dynamically updating interpretations in context, where the falsity of a universal can also occur when there is no *coherent context* in which all of its instances are true.

Whereas S'valuationsim defines classical extensions of partial models, Kamp's contextualism includes contexts in the model. Interpretations of predicates relative to contexts also include positive and negative extensions, but although there might be gaps in an interpretation relative to a context, contexts can be extended to include an interpretation of sentences in the gaps. So, one model, not many. Tolerance holds on Kamp's account. Where $\mathcal{I}_k$ is Kamp's interpretataion function, U is the domain, and $B(c)$ is the set of background assumptions in context c:

(30) $$\mathcal{I}_k(P(x_i))(c) = 1 \text{ iff } \exists a \in U(a \sim \mathcal{I}_k(x_i) \wedge P(a) \in B(c))$$

[13]It should be stressed that although Kamp's exploration into a solution for this problem is very detailed and careful, it is also tentative. Kamp presents a number of possible logical alternatives with respect to adjusting the classical consequence relation, and we will not be able to do the subtleties of his paper justice here.

An object is in the extension of a predicate at a context if the object is tolerantly similar to some object in the extension of the predicate by background assumption. Contexts are dynamic, and so the acceptance of a statement as true modifies the context (which also adds this statement to the background assumptions). Hence sorites series progress via acknowledging tolerance relations. Doing this modifies the base context (thus extending the background assumptions in it). However, as the context is increasingly modified, one may reach a point where a statement is added to the background assumptions that contradicts a statement already in that set. This happens when the acknowledgement of tolerance relations extends far enough to reach the negative extension of the predicate in the base context. The inclusion of a contradiction in the background assumptions means that the (extended) context is *incoherent.*

Hence, although every instance of a tolerance statement $\forall x \forall y((P(x) \wedge x \sim_P y) \rightarrow P(y))$ might be true at a context, if the context is incoherent, the universal statement is nonetheless false. However, unlike supervaluationism, the falsity of the tolerance premise does not entail the truth of a sharp boundary proposition. So:

(31) $$\neg\forall x \forall y((P(x) \wedge x \sim_P y) \rightarrow P(y)) \nvDash_k \exists x \exists y((P(x) \wedge x \sim_P y) \wedge \neg P(y))$$

thus remedying the problem Kamp highlighted with supervaluationism.

Like many other approaches, a form of higher-order problem emerges. As Kamp discusses, it is awkward to explain why ceasing to progress along the sorites series *at a particular moment* is more plausible if created by a switch into an incoherent context, rather than due to a sharp boundary. A similar point is that there may be a point at which, for some a, a is the last true instance of a P in any coherent context. A tentative conclusion is then to see the notion of coherence as itself vague, or coming in degrees. However, that amounts to leaving a new form of vagueness to be explained.

Kamp's contextualism shares with many other analyses to the sorites the commitment that the tolerance premise is false. However, it does seem to assuage this problem in its resolution of an unintuitive result of supervaluationism: Every tolerance conditional *can* be accepted as totally true. However, as Kamp carefully sets about describing, the formal properties of this approach are hard to establish. In particular there are numerous possibilities for how the logical consequence relation could be specified within such as system.

The lasting impact of Kamp's contextualism is apparent in subsequent vagueness research. Kamp introduced into the vagueness literature the idea of dynamically updating contexts which in turn affect the interpretation of propositions containing vague predicates. Manifestations of related ideas are prevalent in subsequent literature. This includes the psychologically described contextualism of Raffman (1994, 1996), the contextualism of Tappenden (1993), van Deemter (1995) and Soames (1999), the dynamic semantics-based approaches of Barker (2002) and Lassiter (2011), and in the similarity sensitivity of the TCS approach (see Section 4.5 and references therein).

4.3.2 The connection between vagueness and context sensitivity

The exact relationship between context sensitivity and vagueness is discussed in various places (Williamson, 1994; Raffman, 1996, 2014; Fara, 2000; Shapiro, 2006, amongst many others). One concern is that context-sensitivity is orthogonal to vagueness. Although it is widely accepted that vague predicates are context-sensitive (the conditions for *tall* said of a mountain differ from *tall* said of a chair), it is disputed whether context-sensitivity is the source of vagueness. Even if we make information about the context highly rich, it does not mean that we will be able to discern where the boundaries in the extensions of predicates lie. We are not freed from vagueness if, for example, we know that *tall* is being applied to a chair, or to a highchair, or to a highchair for 2-3 year olds, since the boundary line for *tall* in each of these contexts is still blurred (and/or *tall* still admits of borderline cases).

One response to this kind of argument is that context switches are dynamic (e.g. Kamp, 1981; Raffman, 1996, 2000) and to some extent arbitrary (Raffman, 1996, 2000; Rayo, 2008). Simply being posed certain questions or shown certain stimuli can evoke a shift in (internal) context. This can mean that, even though there is a (more or less) classical underpinning to our reasoning, the boundaries of vague predicates are elusive in that small changes in context lead them to slip out of our grip. For example, if such context shifts occur even across the assertion of a conjunction, then a contextualist position may be able to explain how one can assert 'P and not-P' without abandoning classical logic or asserting a flat contradiction. If context shifts subtly between the first and second conjunction, then the context to evaluate each conjunct is different. In other words 'P and not-P' could be analysed as asserting 'P (from one perspective), but, at the same time, not-P from another'.

For more discussion of borderline contradictions (not only in relation to (more or less) classical approaches) see, amongst others, Bonini et al. (1999), Alxatib and Pelletier (2011), and Égré and Zehr (2017). For a discussion regarding the kinds of context sensitivity relating to vagueness see, amongst others, Shapiro (2006), Ripley (2011b).

4.4 Epistemically Oriented Theories

4.4.1 Epistemicism

Given the problems generated by departing from classical semantics, the simplest way to avoid such problems is perhaps not to account for vagueness in semantics at all, but rather to explain vagueness as an epistemic phenomenon. This is the proposal that the philosophers Sorensen (1988, 2001) and Williamson (1992, 1994) suggested in a position that became known as *epistemicism*.[14] What we mistake for indeterminacy, is actually ignorance about (or perhaps uncertainty of) the facts.

On epistemicism, no revision to classical semantics based on FOL is necessary. The sorites argument is valid, but unsound because the Tolerance premise is false. The tolerance premise is false because one of its instantiations is false (all the others can be true). Vagueness is a form of ignorance about where this sharp cut-off point is. One might think that epistemicism leaves little work to be done by semanticists working on vagueness since Williamson's and Sorensen's proposals for epistemicism endorse the view that vagueness should leave semantics unaffected (vagueness lies in our epistemic relationship to the meanings of expressions). However, as we shall see in sections 4.4.2 and 4.4.4, other more semantically intricate proposals share with

[14]Since it has influenced some of the positions to be outlined below, we will focus on Williamson's form of epistemicism. Many philosophical subtleties will be overlooked.

epistemicism the conclusion that vagueness is a form of ignorance, or, at least, uncertainty. Hence we shall only briefly outline the epistemicist proposal here.

On Williamson's version of epistemicism, for example, ignorance about sharp boundaries is fleshed out in terms of *margin for error principles* (Williamson, 1994, chs 7-8), namely, that to know that p requires knowing that p with a sufficient margin for error. If our grounds for believing p is too close to a situation in which not p (if it is within a margin for error), then we cannot know p since we would have believed p even if p were false. More specific to the case of vagueness, we know, to some extent, what the extensions of vague predicates are. For example, we probably all know that two metres is tall for an adult human being. However, we do not know exactly what what the extensions of vague predicates are. Our linguistic knowledge is inexact. A person might believe that 186cm is the cut off point for tall, and even if this is true, had the cut-off point been 187cm, she would not have altered her belief in virtue of this difference. Her belief was only true by luck. If a point is too close to the actual boundary (including the boundary point itself), we *cannot* know whether or not that point is the boundary, because our putative knowledge of the meaning of the predicate is too inexact to track any subtle differences in meaning there might have been that would have placed the boundary in a slightly different position.

Nonetheless, a broadly epistemic position is compatible with a range of semantics based on CFOL.[15] We will consider two positions below. One (Fara, 2000; Kennedy, 2007, i.a.) incorporates degrees into the semantics of vague expressions but locates uncertainty in a slightly different place than epistemicism. The other (Lassiter, 2011) enriches dynamic models with Bayesian probability calculus to show how we might be able to track the standards in play for vague expressions without ever being likely to know what these standards are.

With respect to higher order vagueness, since epistemicist accounts reduce vagueness to ignorance about sharp boundaries, a higher order vagueness problem would be, for example, that we would know exactly what would be definitely P, even though definitely P is vague. However, taking Williamson's account and treatment of higher-order vagueness as an example, higher-order vagueness problems can only arise on the basis of endorsing the KK principle (if you know that ϕ, you know that you know that ϕ). Yet there is good reason not to embrace this principle. There are, arguably, things that we know, but that we don't know that we know. Nonetheless, this advantage is won only at the cost of bitting a fairly large bullet with respect to the ignorance we have regarding the meanings of the expressions in our language.

4.4.2 Degree Semantics

Fuzzy approaches reject bivalence and build degrees into semantics. Supervaluationism attempted to resolve some of the resulting problems by removing degrees, but by still rejecting bivalence. Epistemicism does neither, but results in a position which many have found hard to swallow. Another alternative is to incorporate degrees and maintain bivalence. Degree-based semantics has its roots in Bartsch and Vennemann (1972); Cresswell (1977); Bierwisch (1989) (amongst others). Here we will focus on the proposals of Fara (2000) and Kennedy (2007). (In sections 4.4.3 and 4.4.4 two more approaches that adopt this strategy will be discussed.)

[15] Also see MacFarlane (2010) for a proposal for wedding epistemicism with fuzzy logic.

The first departure from the traditional (albeit not first-order) semantics is to add an extra semantic type d (for degree) to the familiar e and t. Gradable adjectives are typed as $\langle e, d\rangle$. Suppose a degree based interpretation function $\mathcal{I}_d$:

(32) $$\mathcal{I}_d(Tall_{\langle e,d\rangle}(a_{\langle e\rangle}))_{\langle d\rangle} = d$$

This expresses that a is tall to degree d. Degree based approaches very successfully capture comparative constructions. Expressions such as 'more' or morphemes such as '-er' are interpreted as functions to inequalities over degrees $\lambda P.\lambda y.\lambda x.F(x) > F(y)$. So a is taller than b ($Tall(a) > Tall(b)$) is true iff the degree of height of a is greater than that of b.

For 'bare' or 'positive' uses of gradable adjectives, a null morpheme *pos* is postulated which has the logical form $\lambda F.\lambda x.F(x) \geq \mathbf{X}_F$, where $\mathbf{X}_F$ is of type $\langle d\rangle$ on a scale determined by the adjective substituted for F. As noted by Fara (2000) and Kennedy (2007), if whatever we supplement for $\mathbf{X}_F$ is not itself context sensitive, then we do not have a very satisfactory analysis of vagueness. Say, for example, that $\mathbf{X}_F$ is interpreted as something like *the average degree to which entities in the relevant class are* F, then the 'cut off' point for F will remain fixed given the contingent facts about entities which are in the extension of F. However, this does not explain why vague predicates admit of borderline cases. We briefly present two solutions to this problem. Interest relativity (Fara, 2000), and domain of discourse restriction (Kennedy, 2007).

Fara (2000) proposes three modifications to the meaning of *pos*: an inclusion of a comparison class property P; $NORM$, a function from a measure function to a function from properties to degrees;[16] and !> a 'significantly greater than' inequality:

(33) $$[pos]_{\text{fara}} = \lambda F.\lambda P.\lambda x.(F(x) \mathrel{!\!>} (NORM(G)(P)))$$

So, given that the relevant property is BB (*being a basketball player*), this yields the proposition for 'a is tall (for a basketball player)' as the following:

(34) $$Tall(a) \mathrel{!\!>} (NORM(Tall)(BB))$$

Where this is interpreted as "a has a significantly greater degree of height than the degree of height that is normal for a basketball player". Crucially, 'significantly greater than' is *interest relative*. What might count as 'significantly greater than' will vary depending on/relative to the agent's interests. The further claim, that yields a treatment of the sorites is that the exact requirements of an agent's interests need not be epistemically accessible to the agent. Hence we are not ignorant of the meanings of vague expressions, but we might be ignorant of exactly what, in a context, the cut off point for being significantly greater than is.

Kennedy (2007) notes an objection to an interest relative account (Stanley, 2003), namely that some bare uses of adjectives seem to be true/false independent of an agent's interests. For example, 'Mount Everest is tall for a mountain', uttered by Jo seems capable of being true irrespective of the interests of the agent involved in asserting it, since the same utterance could have been true even if Jo had never existed (Stanley, 2003, p. 278). Kennedy does not rule out that interest sensitivity may be at play in many cases (Kennedy, 2007, p. 17), but tries

[16]Assuming properties to be of type $\langle e, t\rangle$, this would make $NORM$ of type $\langle\langle e, d\rangle, \langle\langle e, t\rangle, d\rangle\rangle$. Fara does not rule out properties as being interpreted as kinds, however.

to bypass such objections with his context sensitive account on which whether an entity is F is based on distributional, agent insensitive, criteria. Kennedy introduces a context sensitive function, **s**, from measure functions to degrees. The output to this function is a standard of comparison adjective F but sensitive to the context of utterance.

(35) $$[pos]_{\text{kennedy}} = \lambda F.\lambda x.(F(x) \geq \mathbf{s}(F))$$

s selects a degree of F-ness, above which an object 'stands out' as F. The idea is that, unlike the comparative form, the positive form of a gradable adjective is incompatible with "crisp judgements". If a is 187cm in height, and b is 186cm in height, one can felicitously say "a is taller than b", but cannot felicitously say "a is tall compared with b".

For positive form statements, such as "a is tall (for a basketball player)", the function **s** determines a degree of height above which, individuals have a height which, in distributional terms, "stands out" within that class:[17]

(36) $$(Bb(a).Tall(a)) \geq \mathbf{s}(Bb.Tall)$$

We can therefore know the meanings of expressions such as "tall", but what we may be uncertain of is exactly what the contextually determined standard in play is.

This idea, of uncertainty about standards in a context is not analysed much further in Kennedy's account. However, this issue is directly modelled in the Probabilistic Linguistic Knowledge account of Lassiter (2011) to be discussed in section 4.4.4.

4.4.3 Probabilistic Approaches I: Verities

Edgington (1992, 1997) argues that a logic for vagueness shares a structural similarity with classical Bayesian probability calculus. In particular, a logic based on degrees that is similar in structure to probability calculus is in a better position than fuzzy logic to capture logical relations between propositions. Edgington also adopts a range of values $[0, 1]$, however, these are not meant to be degrees of truth, nor are they meant to be degrees of certainty.[18] Instead, they are "degrees of closeness to clear cases of truth" which Edgington dubs as *verities*. In other words, Edgington sees no contradiction in embracing bivalence, but also in assuming a logic for reasoning in cases of vagueness that tracks closeness to clear cases of bivalent values. Verity connectives obey equivalent rules to probabilities. Where $\mathcal{I}_v$ is the verity interpretation function:

Definition 4.2. ***Interpretation of Propositions as Verities.***

1. $\mathcal{I}_v(\phi) \in [0, 1]$
2. $\mathcal{I}_v(\neg\phi) = 1 - \mathcal{I}_v(\phi)$
3. $\mathcal{I}_v(\phi \wedge \psi) = \mathcal{I}_v(\phi) \times \mathcal{I}_v(\psi|\phi)$

[17] Following Heim and Kratzer (1998), "if f is a function of type $\langle\tau, \sigma\rangle$, then $\lambda v : g(v).f(v)$ is a function just like f except that its domain is the subset of things of type τ that satisfy g." (Kennedy, 2007)

[18] However, see Sutton (2015) for a proposal which interprets utterances in a situation theoretic, Bayesian framework based on the probability of a described situation, given some discourse situation.

4. $\mathcal{I}_v(\phi \vee \psi) = \mathcal{I}_v(\phi) + \mathcal{I}_v(\psi) - \mathcal{I}_v(\phi \wedge \psi)$

Edginton's system also has another connective, $|$, whose interpretation is slightly trickier. $\mathcal{I}_v(\psi|\phi)$ should be understood as the degree of closeness to clear truth of ψ, given that ψ is clearly true. Like fuzzy approaches, degrees are written directly into the semantics and interpretations of propositions, however, unlike fuzzy approaches, the connectives $\wedge$ and $\vee$ are not strictly degree-functional, since $\mathcal{I}_v(\phi|\psi)$ is not computable from only $\mathcal{I}_v(\phi)$ and $\mathcal{I}_v(\psi)$. Furthermore, unlike fuzzy systems, verities preserve classical theorems. For example, for all ϕ, $\mathcal{I}_v(\phi \wedge \neg\phi) = 0$ (non-contradiction), and $\mathcal{I}_v(\phi \vee \neg\phi) = 1$ (excluded middle). Furthermore, the difficult cases for fuzzy logic are resolved. If $\mathcal{I}_v(P(a)) = 0.4$ and $\mathcal{I}_v(P(b)) = 0.5$ (b is slightly more clearly P than a), it is still clearly false that a is P, but b is not ($\mathcal{I}_v(P(a) \wedge \neg P(b)) = 0$), since $\mathcal{I}_v(P(a)|\neg P(b)) = 0$.

Notice, however, that, like supervaluationism, a verity-based account also results in a divergence from classical consequence for multi-premise conclusions. Given 4 in Definition 4.2, it is possible that $\mathcal{I}_v(\phi \vee \psi) = 1$ when $\mathcal{I}_v(\phi) < 1$ and $\mathcal{I}_v(\psi) < 1$. To this extent, a verity based approach must answer some of the same criticisms as supervaluationism with respect to this result.

Sorites arguments, on a verity-based account, are analysed as valid but unsound. Every tolerance condition is almost perfectly close to being clearly true, but over repeated steps of *modus ponens*, small degrees of distance from clear truth are introduced. The universal Tolerance premise which is interpreted as a conjunction of its instances is either totally or near completely false, despite having almost completely clearly true instances, and its negation $\exists x \exists y((P(x) \wedge x \sim_P y) \wedge \neg P(y))$ (interpreted as a disjunction of its instances) is either clearly false or very close to clearly false despite all of its instances being almost entirely clearly true. (This differs from supervaluationism in which such existentially quantified statements are true.)

Some philosophical criticisms of Edgington's proposal have been made. For example the interpretation of conditional verities has been questioned (Keefe, 2000). Eddington's account has also been criticised by fuzzy theorists (see, for example, Smith (2008)). The main point of contention is conjunction. Fuzzy accounts face difficulties with logical dependencies between propositions (see section 4.1). However probabilistically grounded accounts face a difficulty with independent propositions. Say that Danny is borderline (0.5) tall and borderline (0.5) old. Fuzzy approaches assign a value of 0.5 to the proposition 'Danny is tall and old'. However, on a Bayesian based calculation, given the independence of the conjunctions, the value is 0.25. Potentially even more worrying is that the decrease of value is exponential (0.5^n) with the number of conjuncts n. For more discussion see Schiffer (2003, §5.4), MacFarlane (2010) Smith (2008, §5.3), Sutton (2013, §8.2).

Despite these criticisms, Edgington's proposal stands as one of the first contemporary applications of Bayesian tools to the problem of vagueness,[19] and has been influential in helping to inspire others to pursue similar probabilisitic avenues.

[19]However see Égré and Barberousse (2014) for a discussion of Emile Borel's account of vagueness, likely the earliest statistical account of vagueness in the 20^{th} century.

4.4.4 Probabilistic Approaches II: Probabilistic Linguistic Knowledge

Barker (2002) develops a non-probabilistic account of vagueness within the dynamic semantics paradigm. He focuses less on the sorites, and more on the impact of uses of vague expressions on the context of discourse. In particular, utterances can carry metalinguistic as well as factual/worldly information. For example: if Ashley is unfamiliar with the standards for tallness in her context, but knows that Billie is 185cm in height, then hearing an utterance of 'Billie is tall' can help Ashley narrow down the standards for tallness in that situation, namely, that the threshold for height must be above 185cm.[20]

Lassiter (2011) builds on this work enriching a Barker-type framework with Bayesian probability calculus so as to be able to model gradience in the interpretations of predicates.[21] Just as there is a good case for being able to represent varying levels of uncertainty in our beliefs about the world, especially as a result of updating our beliefs on the basis of a non-wholly-reliable source, there is also a case for introducing uncertainty about how words are being used.

Lassiter defines a probabilistic belief space that can reflect both worldly and metalinguistic uncertainty. Roughly, worldly uncertainty comes out as a probability distribution over possible worlds. Metalinguistic uncertainty is captured as a probability distribution over precisifications of natural language terms. Formally, Lassiter defines a probabilistic belief space (W, L, μ), and a probability function $\mu : (W, L) \to [0, 1]$ where W is a set of possible worlds and L is a set of possible languages. The probability an agent assigns a possible world will then be the sum of the probabilities of the world-language pairs it occurs in, *mutatis mutandis* for a possible language. Utterances using vague terms will then be interpreted as Bayesian updates on the probabilistic belief space.

Here is an example. For simplicity, assume that our model contains just one possible world (so we have no worldly uncertainty), which is characterised as a proposition about Cam's height.

(37) $$w_1 = \{height(cam) = 188cm\}$$

However, we may be uncertain about what the standard for *tall* (T) is. Assume that our probability space contains five sharp interpretations of tall (T_i):

(38) $$\begin{array}{ll} T_1 = \lambda x.height(x) \geq 150cm; & T_2 = \lambda x.height(x) \geq 160cm; \\ T_3 = \lambda x.height(x) \geq 170cm; & T_4 = \lambda x.height(x) \geq 180cm; \\ T_5 = \lambda x.height(x) \geq 190cm & \end{array}$$

Then μ will be a function that assigns probabilities to world, language pairs. For example:

(39) $$\mu = \{\langle\langle w_1, T_1\rangle, 0.05\rangle, \langle\langle w_1, T_2\rangle, 0.15\rangle, \langle\langle w_1, T_3\rangle, 0.3\rangle, \langle\langle w_1, T_4\rangle, 0.4\rangle, \langle\langle w_1, T_5\rangle, 0.1\rangle\}$$

From this we can calculate the probabilities of w_1 being actual (worldly uncertainty), and of each sharp meaning of tall being the one being used in the context[22] (metalinguistic uncer-

[20] Although itself interesting and influential, we do not further elaborate upon Barker's work here.
[21] Also see Frazee and Beaver (2010). For further developments also see Lassiter and Goodman (2015).
[22] In this case, both are pretty trivial given that we have only one possible world.

tainty):

(40) $$\mu(w_1) = 1$$

(41) $$\mu(T_1) = 0.05,\ \mu(T_2) = 0.15,\ \mu(T_3) = 0.3,\ \mu(T_4) = 0.4,\ \mu(T_5) = 0.1$$

The probability that 'Cam is tall' is true, is then calculated as the sum of the probabilities of the world, language pairs in which it is true that Cam is tall, weighted against the probability that the world in the pair is the actual world. There is only one world language pair in which it is false that Cam, at 188cm in height, is tall, namely $\langle w_1, T_5\rangle$. This gives the probability of the truth of 'Cam is tall':

(42) $$\begin{aligned} p(\text{Cam is tall}) &= 0.05 \times 1 + 0.15 \times 1 + 0.3 \times 1 + 0.4 \times 1 \\ &= 0.9 \end{aligned}$$

As with the other epistemic approaches, there are on Lassiter's approach, sharp boundaries for vague predicates. At least, however, on this probabilistic variant, we regain a notion of gradience (in terms of slowly increasing/decreasing probabilities of the truth of propositions as the sorites series progresses). Furthermore, via the dynamic aspects of this account, we also have a more systematic explanation of how agents can keep track of shifting standards of interpretation in context.

A suggestion for a treatment of higher order vagueness is given in Lassiter (2011). NL expressions are interpreted as distributions over possible precise meanings. Expressions such as 'clearly tall' can also be interpreted as such. We may know that the threshold for 'clearly tall' is higher that the threshold for 'tall', but there is no reason to think that we should be any more certain where the former threshold lies. Furthermore, although not explicitly mentioned by Lassiter, a more detailed approach could be derived as a fairly obvious extension of Lassiter's work on epistemic adverbs such as *certainly* and *possibly* which also includes a representation of uncertainty about the evidence for which one makes an assertion (Lassiter, 2016).

4.5 Tolerant, Classical, Strict

The final proposed solution to the puzzles of vague language that we will discuss is the *Tolerant, Classical, Strict* solution (Cobreros et al., 2012). This system is one of a series of recent logical systems (such as Frankowski, 2004; Zardini, 2008; Smith, 2008, among others) that explore the use of a *permissive consequence* relation in the resolution of semantic paradoxes such as the sorites and the liar. In other words, in TCS and other systems like it, rather than expressing the relation of the preservation of a distinguished truth value, the consequence relation allows a 'weakening' of standards when going from premises to conclusion.

The TCS system was originally developed as a way to allow vague predicates to be tolerant (that is, to satisfy $\forall x \forall y[P(x)\ \&\ x \sim_P y \rightarrow P(y)]$, without running into the sorites paradox. The paradox is avoided in this system through adopting a consequence relation that has different properties from the one found in FOL. More specifically, TCS departs from classical logic in that it adopts three notions of satisfaction: classical truth, tolerant truth, and its dual, strict truth. As such, the system violates the principle of bivalence. Formulas are tolerantly/strictly satisfied based on classical truth and predicate-relative, possibly non-transitive *indifference relations*, which are encoded in the model. For a given predicate P, an indifference relation, $\sim_P$, relates those individuals that are viewed as sufficiently similar with respect to P. For example, for the predicate *tall*, $\sim_{tall}$ would be something like the relation "not looking to have distinct heights".

Formulas in TCS are interpreted into T(olerant) models, defined as follows:

Definition 4.3. *T(olerant) Model. A t-model is a tuple $\langle D, m, \sim \rangle$, where $\langle D, m \rangle$ is a model (as defined in section 2) and $\sim$ is a function that takes any predicate P to a binary relation $\sim_P$ on D. For any P, $\sim_P$ is reflexive and symmetric (but possibly not transitive).*

In TCS, classical satisfaction (written $\vDash^c$) is just that: it has all the properties that we discussed in section 2. Therefore, the classical interpretation function $\mathcal{I}$ is total. Additionally, in order to be able to refer to indifference relations in formulas, the TCS language has, for every predicate P, a binary predicate I_P, which is classically interpreted as denoting the indifference relations associated with P (i.e. for terms t_1, t_2, $\mathcal{I} \vDash^c t_1 I_P t_2$ iff $\mathcal{I}(t_1) \sim_P \mathcal{I}(t_2)$).

Tolerant and strict satisfaction ($\vDash^t$ and $\vDash^s$) are defined based on classical satisfaction and the $\sim$ relations. Informally, in this framework, we say that *John is tall* is tolerantly true just in case John has a very similar height to someone who is classically tall (i.e. has a height greater than or equal to the contextually given 'tallness' threshold). Likewise, we say that *John is tall* is strictly true just in case everyone whose height is similar to John's is classically tall. Formally, the definitions are given as follows for atomic fomulas, formulas with indifference relations or negation, the conditional and the universal quantifier. The definitions of satisfaction for other connectives and other quantifiers are straightforward:

Definition 4.4. ***Tolerant/Strict satisfaction***($\vDash^{t/s}$)*. Let $\mathcal{I}$ be an interpretation. For all predicates P and terms t_1, t_2:*

1. $\mathcal{I} \vDash^t P(a_1)$ *iff* $\exists a_2 \sim_P a_1 : \mathcal{I} \vDash^c P(a_2)$
2. $\mathcal{I} \vDash^t t_1 I_P t_2$ *iff* $\mathcal{I}(t_1) \sim_P \mathcal{I}(t_2)$
3. $\mathcal{I} \vDash^t \neg\phi$ *iff* $\mathcal{I} \nvDash^s \phi$
4. $\mathcal{I} \vDash^t \phi \rightarrow \psi$ *iff if* $\mathcal{I} \nvDash^s \phi$ *or* $\mathcal{I} \vDash^t \psi$
5. $\mathcal{I} \vDash^t \forall x_1 \phi$ *iff for every* a_1 *in* D, $\mathcal{I}[a_1/x_1] \vDash^t \phi$
6. $\mathcal{I} \vDash^s P(a_1)$ *iff* $\forall a_2 \sim_P a_1 : \mathcal{I} \vDash^c P(a_2)$
7. $\mathcal{I} \vDash^s t_1 I_P t_2$ *iff* $\mathcal{I}(t_1) \sim_P \mathcal{I}(t_2)$
8. $\mathcal{I} \vDash^s \neg\phi$ *iff* $\mathcal{I} \nvDash^t \phi$
9. $\mathcal{I} \vDash^s \phi \rightarrow \psi$ *iff if* $\mathcal{I} \nvDash^t \phi$ *or* $\mathcal{I} \vDash^s \psi$
10. $\mathcal{I} \vDash^s \forall x_1 \phi$ *iff for every* a_1 *in* D, $\mathcal{I}[a_1/x_1] \vDash^s \phi$

Note that the predicates that refer to indifference relations are interpreted 'crisply' (in the words of Cobreros et al. (2012)): their interpretation is the same on all kinds of satisfaction.

The framework has three notions of satisfaction, and from these notions we can derive 9 consequence relations (defined in a similar manner to the consequence relation of FOL in definition 2.8). As discussed in Cobreros et al. (2012), these relations are in the following lattice order (based on inclusion), where $\vDash^{mn}$ stands for reasoning from m interpreted premises to n interpreted conclusions. Note (as shown in Cobreros et al. (2012)) that $\vDash^{cc}$ is equivalent to consequence in classical FOL (i.e. reasoning from classical premises to classical conclusions). Furthermore, $\vDash^{tt}$ is equivalent to consequence in Priest (1979)'s *Logic of Paradox* (LP), and $\vDash^{ss}$ is equivalent to strong Kleene logic (K3).

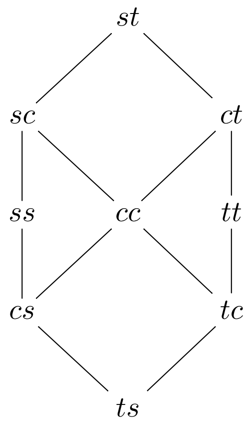

Figure 1 – Consequence relations in TCS

In their article, Cobreros et al. (2012) argue that the system that is the most appropriate for modelling natural language reasoning is the $\Vdash^{st}$ system; that is, reasoning from strictly interpreted premises to tolerantly interpreted ones. As such, TCS (with $\vDash^{st}$) explains the puzzling properties of vague language in the following way: Firstly, although classical negation partitions the domain (like it does in FOL), the definition of tolerant negation actually allows for $P(a_1)$ and $\neg P(a_1)$ to be tolerantly true for some individual a_1. Individuals like a_1 are the borderline cases. The reason that we have difficulty deciding whether a borderline individual is part of a predicate's extension or anti-extension is that such an individual is actually part of both sets. In other words, at the level of tolerant truth, TCS is paraconsistent: contradictions involving borderline cases do not result in explosion (like they do in classical logic). Secondly, TCS preserves the intuition behind the fuzzy boundaries/tolerance property because the principle of tolerance is, in fact, valid at the level of tolerant truth. Note that it is neither classically valid nor strictly valid.

How this system avoids the sorites paradox is a bit more complicated. Firstly, following (Cobreros et al., 2012, 27), we can distinguish two syntactic versions of the argument. The first version proceeds directly from indifference relations:

(43) **Sorites version 1:**

a. $P(a_1)$
b. $\forall i \in [1, n](a_i I a_{i+1})$

---

c. $P(a_k)$

This version of the Sorites is *st*-invalid. However, what is interesting is that TCS (with $\vDash^{st}$) validates each step along the way, which seems appropriate.

(44) Step-wise Tolerance
a. $P(a_1)$
b. $a_1 I a_2$

c. $P(a_2)$

The reason that (43) is invalid, despite the validity of (44) for all individuals adjacent on the scale, is that $\vDash^{st}$ is not transitive.

There is, however, a second version of the Sorites which is more similar to the formulation presented in (13) and *is st*-valid:

(45) **Sorites version 2:**
a. $P(a_1)$
b. $\forall i \in [1, n](a_i I_P a_{i+1})$
c. $\forall x \forall y((P(x) \wedge x I_P y) \rightarrow P(y))$

d. $P(a_k)$

However, we still avoid paradox. Although (45) is valid, it is not sound. Recall that, with $\vDash^{st}$, we are reasoning from strict premises to tolerant conclusions. As mentioned above, the principle of tolerance is neither c-valid nor s-valid; thus, (45-c) will never be strictly true.

Although, with its non-classical consequence relation, the TCS system allows for the (tolerant) validity of the tolerance principle, an important open question in this framework is the treatment of higher order vagueness. Unlike some of the approaches discussed above, TCS solution to the first-order sorites does not immediately extend to possible higher order sorites created by the introduction of operators like *clearly/definitely/determinately.* However, as Cobreros et al. (forthcoming) point out, whether this is such a big problem is not so clear. Since something like a 'determinately/clearly' operator is already implicitly built into the definition of strict satisfaction, it turns out that properly formulating the higher-order vagueness paradox within TCS (with $\vDash^{st}$) is a bit trickier than for some other theories. This is because the *clearly/definitely* operators themselves would have to be given strict and tolerant interpretations, and so there may be room for avoiding a higher order paradox through defining the tolerant/strict interpretations of operators in an appropriate way or even through prohibiting them altogether. This being said, even if we set the question of determinateness operators aside, within the TCS system, borderline cases still do have sharp boundaries, so the treatment of 'borderline cases of borderline cases' is left open at this point[23].

[23] For example, (Cobreros et al., forthcoming, 18) say, "Whether we can accommodate indefinitely iterated borderline cases is an issue that goes beyond the scope of this paper, but our main point, once more, is that the semantics of determinateness is a matter distinct from strict assertion proper." (Ripley, 2013a, §4) gives an account of HOV in S'valuationism and LP, which would probably be applicable to TCS; however, the exploration of this idea is out of the scope of this review paper.

4.6 Summary

Table 1 includes a summary of the key semantic differences between the approaches just discussed. Some details are not mentioned however, such as the subtle ways in which the consequence relation differs between different accounts.

Classical FOL theorems and definitions	Fuzzy Logic	Super-valuationism	Sub-valuationism	Kamp's Contextualism	Epistemicism, Degree Sem. Verities PLK	Tolerant, Classical Strict $\vDash^{st}$
$\neg\forall x\phi \vDash \exists x.\neg\phi$	✓	✓	✓	✗	✓	✓
Total $\mathcal{I}$ function	✓	✗	✗	✓	✓	✓
No gaps	✗	✗	✓	✓	✓	$✗_s$ $✓_t$
No gluts	✓	✓	✗	✓	✓	$✓_s$ $✗_t$
Excluded Middle	✗	✓	✓	✓	✓	$✗_s$ $✓_t$
Non-contradiction	✗	✓	✓	✓	✓	$✓_s$ $✗_t$

Table 1 – Logico-Semantic Commitments of Theories of Vagueness

5 Conclusion

Research on vagueness over the last forty years has become both increasingly inter-disciplinary and increasingly vast. Here we have given an overview of what we take to be some of the most important lines of research for linguists and semanticists. To conclude this overview, we wish to make a few observations about past and present research, but also speculate a little on the future.

From a theoretical-historical perspective, vagueness research has, up until recently, been a story of increasing semantic conservatism. Looking back to Table 1, one can notice a slow move towards classicism from fuzzy logical approaches, to supervaluationism, to contextualism, to epistemic theories. However, far from it being a lesson learned by semanticists to avoid even trying to model vagueness, the epistemicist attitude towards the impact of vagueness on logic and semantics has not become dominant within linguistics. Granted, many descendant positions of epistemicism do employ epistemic explanations with respect to the presence of sharp boundaries at some point in their respective theories, however it is notable how even these theories differ from epistemicism with respect to the need for semantic innovation as a reaction the phenomena of vagueness. We suspect that the move away from epistemicism within linguistics has principally been motivated by some of the empirical factors we have mentioned.

First, the impact of empirical considerations can be seen in the examination of the difference in scale structure between relative and absolute adjectives which has an empirical grounding in the interaction between adjectives and absolute intensifiers (*completely/totally #tall/straight*). The enrichment of semantics with degrees has also been the result of trying to accommodate comparative constructions, and especially the seeming non-vagueness of comparative forms compared with positive forms (*taller than* is not (as vague as) *tall* in its positive form).

Second, attention to communication and the impact of context has led to some departures from epistemicism. A good example of this is Lassiter (2011), who criticises (Williamson's version of) epistemicism for the implausibility of divorcing meaning from humans' knowledge of its use: "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 [Williamson's] epistemicism tends to come across as a *reductio ad absurdum* of the theory." (Lassiter, 2011, p. 128). Barker and Lassiter's innovations were intended to begin to get a handle on how the way an expression is used in a situation can give one metalinguistic information with which to narrow down the possible interpretations available for that expression. Relatedly, vagueness could be seen as being a more pragmatic, or a more semantic phenomenon. In the former case, many of the accounts discussed could be used as formal characterisations of the use conditions for vague predicates (where the truth conditions could remain classical).

Third, the empirical data mentioned in section 3.1 has led to new developments too. An empirically based theory of vagueness cannot ignore the large amount of evidence that people typically use either F *and not* F or *neither* F *nor not* F constructions when faced with judging borderline cases using a vague predicate F. Hence some more recent departures from classicism (such as subvaluationism and TCS) have the express purpose of being able to interpret such utterances literally without deeming the speakers helplessly irrational. However, there are some classical avenues for addressing these data such as forms of contextualism.

Looking to the present, a large number of questions regarding the modelling of vagueness remain open to debate. For example, there is still no consensus as to whether vagueness is, at root, a semantic or an epistemic/doxastic phenomenon. At one point in time, the question of higher-order vagueness looked to lean in favour of more epistemically oriented accounts, however, as we saw in section 4.5 there are prospects for more semantically oriented approaches to account for higher orders of vagueness, too.

Looking to the future, although we are sceptical that any one account of vagueness will emerge any time soon as *the* answer to all of the subtle facts to be accommodated, we are also optimistic about the direction of vagueness research. As we have seen, it is still widely held to be a desideratum of a theory of vagueness that it provides a non-paradoxical outcome for sorites arguments, however, as we have just discussed, current and more recent accounts of vagueness increasingly aim at covering much broader sets of syntactic, semantic, and pragmatic data. If there is, as there seems to be, a rich source of such data, both discussed and yet to be discussed in the literature, the study of vagueness probably has a long future within linguistic and semantic theory.

6 Further Reading

We have not been able to cover all of the topics relating to vagueness in semantics. Most notably, we have not been able to address vagueness in lexical semantics. For further discussion of the issues discussed in this chapter, and for discussion of some of the issues not discussed here, see chapters in Handbooks such as van Rooij (2011) and Kamp and Sassoon (2016), and volumes of collected papers on vagueness such as Dietz and Moruzzi (2009), Égré and Klinedinst (2011), Nouwen et al. (2011), and Cintula et al. (2011).

References

Sam Alxatib and Jeffery Pelletier. The psychology of vagueness: Borderline cases and contradictions. *Mind & Language*, 26(3):287–326, 2011.

Sam Alxatib, Peter Pagin, and Uli Sauerland. Acceptable contradictions: Pragmatics or semantics? a reply to cobreros et al. *Journal of Philosophical Logic*, 42:619–634, 2013.

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

Renate Bartsch and Theo Vennemann. The grammar of relative adjectives and of comparison. *Linguistische Berichte*, 20:19–32, 1972.

Manfred Bierwisch. The semantics of gradation. In Manfred Bierwisch and E Lang, editors, *Dimensional Adjectives*, pages 71–261. Springer, Berlin, 1989.

Susanne Bobzien. I—columnar higher-order vagueness, or vagueness is higher-order vagueness. *Aristotelian Society Supplementary Volume*, 89(1):61–87, 2015.

Nicalao Bonini, Daniel Osherson, Riccardo Viale, and Timothy Williamson. On the psychology of vague predicates. *Mind & Language*, 14:377–393, 1999.

Émile Borel. An economic paradox: The sophism of the heap of wheat and statistical truths. *Erkenntnis*, 79(5):1081–1088, 1907/2014.

Heather Burnett. *The Grammar of Tolerance: On Vagueness, Context-Sensitivity, and the Origin of Scale Structure.* PhD thesis, University of California, Los Angeles, 2012.

Heather Burnett. A Delineation solution to the puzzles of absolute adjectives. *Linguistics & Philosophy*, 37:1–39, 2014.

Gennaro Chierchia and Sally McConnell-Ginet. *Meaning and grammar: an introduction to semantics.* MIT Press, Cambridge, 2000.

Petr Cintula, Christian G. Fermüller, Lluís Godo, and Petr Hájek, editors. *Understanding Vagueness: Logical, Philosophical, and Linguistic Perspectives.* College Publications, 2011.

Pablo Cobreros. Paraconsistent vagueness: A positive argument. *Synthese*, 183(2):211–227, 2011.

Pablo Cobreros, Paul Égré, David Ripley, and Robert van Rooij. Tolerance and mixed consequence in the s'valuationist setting. *Studia Logica*, page ??, 2011.

Pablo Cobreros, Paul Égré, David Ripley, and Robert van Rooij. Tolerant, classical, strict. *Journal of Philosophical Logic*, 41:347–385, 2012.

Pablo Cobreros, Paul Egré, David Ripley, and Robert van Rooij. Vagueness, truth and permissive consequence. In Theodora Achourioti, Henri Galinon, José MartÃŋnez Fernández, and Kentaro Fujimoto, editors, *Unifying the Philosophy of Truth.* Springer, forthcoming.

Max Cresswell. The semantics of degree. In Barbara Partee, editor, *Montague Grammar*, pages 261–292. Academic Press, New York, 1977.

DA Cruse. *Lexical Semantics.* Cambridge University Press, Cambridge, UK, 1986.

Richard Dietz and Sebastiano Moruzzi, editors. *Cuts and Clouds. Vaguenesss, its Nature and its Logic.* Oxford University Press, 2009.

Dorothy Edgington. Validity, uncertainty and vagueness. *Analysis*, 52(4):193–204, 1992.

Dorothy Edgington. Vagueness by degrees. In R. Keefe and P. Smith, editors, *Vagueness: A Reader*, pages 294–316. MIT Press, Cambridge, MA, 1997.

Paul Égré and Anouk Barberousse. Borel on the heap. *Erkenntnis*, 79(5):1043–1079, 2014.

Paul Égré and Denis Bonnay. Vagueness, uncertainty, and degrees of clarity. *synthese*, 154:–, 2010.

Paul Égré and Nathan Klinedinst, editors. *Vagueness and Language Use.* Palgrave Macmillan, 2011.

Paul Égré and Jérémy Zehr. Are gaps preferred to gluts? a closer look at borderline contradictions. In E. Castroviejo, G. Weidman Sassoon, and L. McNally., editors, *The Semantics of Gradability, Vagueness, and Scale Structure - Experimental Perspectives.* Springer, 2017.

Paul Egré, Vincent De Gardelle, and David Ripley. Vagueness and order effects in color categorization. *Journal of Logic, Language and Information*, 22(4):391–420, 2013.

Matti Eklund. What vagueness consists in. *Philosophical Studies*, 125:27–60, 2005.

Delia Graff Fara. Shifting sands: An interest-relative theory of vagueness. *Philosophical Topics*, 28, 2000.

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

S Frankowski. Formalization of a plausible inference. *Bulletin of the Section of Logic*, 33: 41–52, 2004.

Joey Frazee and David Beaver. Vagueness is rational under uncertainty. In Maria Aloni, Harald Bastiaanse, Tikitu de Jager, and Katrin Schulz, editors, *Logic, Language and Meaning: 17th Amsterdam Colloquium, Amsterdam, The Netherlands, December 16-18, 2009, Revised Selected Papers, Lecture Notes in Artificial Intelligence*, pages 153–162. Springer, 2010.

J. A. Goguen. The logic of inexact concepts. *Synthese*, 19(3-4):325–373, 1969.

Irene Heim and Angelika Kratzer. *Semantics in Generative Grammar.* Cambridge: Blackwell, 1998.

Dominic Hyde. From heaps and gaps to heaps of gluts. *Mind*, 106:641–660, 1997.

Dominic Hyde. *Vagueness, Logic and Ontology.* Ashgate, 2008.

Dominic Hyde and Mark Colyvan. Paraconsistent vagueness: Why not? *Australasian Journal of Logic*, 6:107–121, 2008.

Hans Kamp. Two theories about adjectives. In Edward Keenan, editor, *Formal Semantics of Natural Language*, pages –. Cambridge University Press, Cambridge, 1975.

Hans Kamp. The paradox of the heap. In Uwe Mönnich, editor, *Aspects of philosophical logic.* Riedel, Dortrecht, 1981.

Hans Kamp and Barbara Partee. Prototype theory and compositionality. *Cognition*, 57: 129–191, 1995.

Hans Kamp and Antje Rossdeutscher. DRS-construction and lexically driven inferences. *Theoretical Linguistics*, 20:165–235, 1994.

Hans Kamp and Galit Sassoon. Vagueness. In Maria Aloni and Paul Dekker, editors, *The Cambridge Handbook of Formal Semantics*. Cambridge University Press, 2016.

Rosanna Keefe. *Theories of vagueness*. Cambridge University Press, Cambridge, 2000.

Edward Keenan and Leonard Faltz. *Boolean Semantics for Natural Language*. Reidel, Dordrecht, 1985.

Christopher Kennedy. Vagueness and grammar: The study of relative and absolute gradable predicates. *Linguistics and Philosophy*, 30:1–45, 2007.

Christopher Kennedy and Louise McNally. Scale structure and the semantic typology of gradable predicates. *Language*, 81:345–381, 2005.

Manfred Krifka. Nominal reference, temporal constitution and thematic relations. In Anna Szabolcsi and Ivan Sag, editors, *Lexical Matters*, pages 29–53. CSLI Publications, Stanford, 1989.

Daniel Lassiter. Vagueness as probabilistic linguistic knowledge. In R. Nouwen, U. Sauerland, H.C. Schmitz, and R. van Rooij, editors, *Vagueness in Communication*. Springer, 2011.

Daniel Lassiter. Must, knowledge, and (in)directness. *Natural Language Semantics*, 24(2): 117–163, 2016. ISSN 1572-865X. doi: 10.1007/s11050-016-9121-8. URL http://dx.doi.org/10.1007/s11050-016-9121-8.

Daniel Lassiter and Noah Goodman. Adjectival vagueness in a bayesian model of interpretation. *Synthese*, pages 1–36, 2015.

David Lewis. General semantics. *Synthese*, 22:18–67, 1970.

Godehard Link. The logical analysis of plurals and mass nouns: A lattice-theoretic approach. In Rainer Bauerle, Cristophe Schwartze, and Arnim von Stechow, editors, *Meaning, Use and the interpretation of language*, pages 302–322. Mouton de Gruyter, The Hague, 1983.

Jan Łukasiewicz. A numerical interpretation of the theory of propositions. In L. Borkowski, editor, *Selected Works*, pages 129–130. North Holland, 1922/1970.

John MacFarlane. Fuzzy epistemicism. In Richard Dietz and Sebastiano Moruzzi, editors, *Cuts and Clouds. Vaguenesss, its Nature and its Logic*. Oxford University Press, 2010.

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

R. Nouwen, U. Sauerland, H.C. Schmitz, and R. van Rooij, editors. *Vagueness in Communication*. Springer, 2011.

Charles Sanders Peirce. Vague. In JM Baldwin, editor, *Dictionary of Philosophy and Psychology*, page 748. Macmillan, New York, 1902.

Manfred Pinkal. *Logic and Lexicon*. Kluwer Academic Publishers, Dordrecht, 1995.

Graham Priest. Logic of paradox. *Journal of Philosophical Logic*, 8:219–241, 1979.

Diana Raffman. Vagueness without paradox. *Philosophical Review*, 103(1):41–74, 1994.

Diana Raffman. Vagueness and context-relativity. *Philosophical Studies*, Vol. 81:175–192, 1996.

Diana Raffman. Is perceptual indiscriminability nontransitive? *Philosophical Topics*, 28: 153–175, 2000.

Diana Raffman. *Unruly words: a study of vague language.* OUP, Oxford, 2014.

Agustín Rayo. Vague representation. *Mind*, 117(466):329–373, 2008.

David Ripley. Contradictions at the borders. In Rick Nouwen, , Robert van Rooij, Uli Sauerland, and Hans-Christian Schmitz, editors, *Vagueness in Communication*, page forthcoming. Springer, 2011a.

David Ripley. Inconstancy and inconsistency. In Petr Cintula, Christian G. Fermüller, Lluís Godo, and Petr Hájek, editors, *Understanding Vagueness: Logical, Philosophical, and Linguistic Perspectives*, pages 41–58. College Publications, 2011b.

David Ripley. Sorting out the sorites. In *Paraconsistency: Logic and applications*, pages 329–348. Springer Netherlands, 2013a.

David Ripley. Sorting out the sorites. In Koji Tanaka, Francesco Berto, Edwin Mares, and Francesco Paoli, editors, *Paraconsistency: Logic and Applications*, pages 329–348. Springer, 2013b.

Stephen R. Schiffer. *The Things We Mean.* Oxford University Press, 2003.

Phil Serchuk, Ian Hargreaves, and Richard Zach. Vagueness, logic and use: Four experimental studies on vagueness. *Mind & Language*, 26(5):540–573, 2011.

Stewart Shapiro. *Vagueness in Context.* Oxford University Press, Oxford, 2006.

Nicolas Smith. *Vagueness and degrees of truth.* Oxford University Press, Oxford, 2008.

Scott Soames. *Understanding Truth.* Oxford University Press, New York, 1999.

Roy Sorensen. *Blindspots.* Clarendon Press, Oxford, 1988.

Roy Sorensen. *Vagueness and Contradiction.* Clarendon Press, Oxford, 2001.

Jason Stanley. Context, interest relativity and the sorites. *Analysis*, 63(4):269–281, 2003.

Peter. R. Sutton. *Vagueness, Communication, and Semantic Information.* PhD thesis, King's College London, 2013.

Peter R. Sutton. Towards a probabilistic semantics for vague adjectives. In Hans-Christian Schmitz and Henk Zeevat, editors, *Language, Cognition, and Mind.* Springer, 2015.

Peter R. Sutton. Probabilistic approaches to vagueness and semantic competency. *Erkenntnis*, Jun 2017.

Jamie Tappenden. The liar and sorites paradoxes: towards a unified treatment. *Journal of Philosophy*, 90:551–577, 1993.

Assaf Toledo and Galit Sassoon. Absolute vs relative adjectives: Variation within or between individuals. In *Proceedings of Semantics and Linguistic Theory 21*, page to appear, 2011.

Kees van Deemter. The sorites fallacy and the context-dependence of vague predicates. In Makoto Kanazawa, Christopher Pinon, and Henriette de Swart, editors, *Quantifiers, Deduction, and context*, pages 59–86. CSLI Publications, Stanford, 1995.

Robert van Rooij. Vagueness, tolerance and non-transitive entailment. Unpublished manuscript, University of Amsterdam, 2010.

Robert van Rooij. Vagueness and linguistics. In G Ronzitti, editor, *The vagueness handbook*, page forthcoming. Springer, Dordrecht, 2011.

J. Robert G. Williams. Degree supervaluational logic. *The Review of Symbolic Logic*, 4: 130–149, 2011.

Timothy Williamson. Vagueness and ignorance. *Proceedings of the Aristotelian Society*, 66: 145–162, 1992.

Timothy Williamson. *Vagueness*. Routledge, London, 1994.

Yoad Winter. *Flexibility Principles in Boolean Semantics*. MIT Press, Cambridge, 2001.

Crispin Wright. On the coherence of vague predicates. *Synthese*, 30:325–365, 1975.

Crispin Wright. The illusion of higher-order vagueness. In Richard Dietz and Sebastiano Moruzzi, editors, *Cuts and Clouds. Vaguenesss, its Nature and its Logic*, pages 523–549. Oxford University Press, 2009.

Y Yoon. Total and partial predicates and the weak and strong interpretations. *Natural Language Semantics*, 4:217–236, 1996.

Lotfi Zadeh. Fuzzy sets. *Information and Control*, 19(8):328–353, 1965.

Elia Zardini. A model of tolerance. *Studia logica*, 90:337–368, 2008.