

# Vagueness and Precifications

Mats Grimsgaard



Master's Thesis in Philosophy

Supervisor: Olav Asheim

Department of Philosophy, Classics, History of Art and Ideas

UNIVERSITY OF OSLO

Spring 2014



# *Abstract*

The aim of this thesis is to explore the role of precisifications of vague predicates. Basically, the idea is that vague predicates can be, and are in fact, made more precise without altering their underlying concepts or truth-conditions. Vague terms and their precisifications are conceived as more or less precise instances of the same lexical entities. Precisifications are employed in several leading views on how we interpret and understand vague language, and how we utilize and reason with partial concepts and inexact knowledge—which covers most, if not all, human thinking and speaking.

Although precisifications play important roles in the semantics of vagueness, they are not as straightforwardly understood as they might appear. There is more than one way to cash out the notion of precisifications, which have important implications for what extent various theories might be seen as a sufficient analysis of vagueness itself, as opposed to mere simulations of vagueness. Yet, many authors seem to have little concern for this issue. As a result, a term that is central to some of the most popular responses to vagueness might turn out to be ambiguous, if not vague.

In chapter 1 I present and discuss the most vicious feature of vagueness: its tendency to generate paradoxes. This provides a general overview of the topic and introduces some important terms and concepts. Chapter 2 is a discussion of how to cash out the notion of precisifications, and not least ‘admissible precisifications’. I show that the notion is tied to the extensions of predicates. While extensions can be provided in terms of set theory, the notion of admissibility remains murky; beyond a certain point we are unable to separate the ‘good’ from the ‘bad’ precisifications. In chapter 3 I discuss Gottlob Frege’s so-called sharpness requirement, and how it works to protect formal languages from vagueness.

---

We see how Frege's conception of extensions as sets causes problems, and I question whether extensions are apt to characterize predicates at all. Finally we look at some of Frege's often overlooked views on ordinary language. In chapter 4 we consider a problem with the soritical arguments, which gives a fresh perspective on precisifications. I argue that precisifications provide no explanation of soriticality, and close off by briefly discussing the element of risk in theories of the semantics of vagueness.

# Contents

<b>Introduction</b>	<b>v</b>
<b>1 The Sorites Paradox</b>	<b>1</b>
1.1 Various forms of the sorites paradox . . . . .	1
1.2 Strong and weak tolerance . . . . .	4
1.3 Margin of error . . . . .	7
1.4 Nihilism . . . . .	11
1.5 Supervaluationism . . . . .	14
1.6 Opposing views, you're both right . . . . .	16
<b>2 Precifications of Vague Predicates</b>	<b>19</b>
2.1 Vague and precise predicates . . . . .	19
2.2 Admissible precisifications . . . . .	23
2.3 The relation between a vague predicate and its precisifications . . . . .	29
2.4 Centipede games and unsharpenable vagueness . . . . .	32
<b>3 Frege on Sharpness</b>	<b>39</b>
3.1 The sharpness requirement . . . . .	39
3.2 Axiom V . . . . .	44
3.3 Frege on ordinary language . . . . .	50
<b>4 The Many-boundaries Approach</b>	<b>55</b>
4.1 The forced march phenomena . . . . .	55
4.2 Forced marches reveal nothing interesting about vagueness . . . . .	59
4.3 Scrambled sorites . . . . .	63
4.4 Quantities . . . . .	67

CONTENTS

---

4.5 Observation and risk . . . . .	70
<b>Bibliography</b>	<b>73</b>

# *Introduction*

## VAGUENESS

The only barrier to enumerability is to be found in the imperfection of concepts. Bald people for example cannot be enumerated as long as the concept of baldness is not defined so precisely that for any individual there can be no doubt whether he falls under it.<sup>1</sup>

Frege, Letter to Marty, 1882

There is an embarrassing fact that we cannot account for the precise meaning of most words we use in ordinary language. We think we use words like ‘bald’, ‘tall’, ‘yellow’, and so on, to make true statements about the world. Yet, we cannot specify exactly how bald a bald man must be, or precisely when a green banana turns yellow. Apparently, these words have no precise boundaries of application; a banana does not turn yellow in an instant. It is a slow and gradual process that involves many stages in which we may not be able to determine whether the banana is green or yellow. That is vagueness.

Vagueness is important. It involves issues that are far more serious than the philosopher’s petty games of determining the colour of a banana, or whether your neighbour is bald. Words like ‘person’, ‘life’, ‘child’, ‘freedom’ and so on, are vague too. Slogans in ethics like ‘Life has intrinsic value’ has no determinate meaning, as we clearly see in debates about abortion and euthanasia; most of these debates do not turn on whether life has intrinsic value or whether a person has a right to live, but rather on how we should draw boundaries for the application of words like ‘life’ and ‘person’.

---

<sup>1</sup>Gottfried Gabriel et al. 1980, p. 163.

The question ‘What is vagueness?’ has no straightforward answer. There are many rival views on vagueness, concerning its nature, its range, and its resolution. Fundamental questions about vagueness are still unresolved. But when we restrict our scope to vagueness in language and predicates, three features seems to jointly characterize the vagueness phenomenon:<sup>2</sup>

*Borderline cases.* Sometimes it cannot be determined whether a predicate applies. For instance whether a given banana is yellow, or whether an adolescent is a child. These are borderline cases.

*Tolerance.* Vague predicates are tolerant, in the sense that they are insensitive to tiny variations. For instance a man does not turn bald on the loss of one single hair (regardless of his initial state).

*Sorites-susceptibility.* Vague predicates are susceptible to Sorites paradoxes. If you have a heap of sand, then the net removal of one single grain does not destroy the heap. (‘Heap’ is vague, by the way.) But if you keep removing the grains one by one, the heap will eventually disappear. Nevertheless, no single grain-removal is accountable for destroying the heap, so according to our logic there should still be a heap left.

These features pose problems on different levels. The problem with borderline cases seems to be primarily epistemic; we do not know whether an almost ripe banana is yellow, and it seems to be no way of knowing. But this problem is only epistemic if we suppose that there are unknowable facts about the yellowness of almost ripe bananas. We may deny that there are such facts; the almost ripe banana is neither yellow nor non-yellow. This yields a truth-value gap in the predicate ‘yellow’, which poses a problem for classical logic, that expects every

---

<sup>2</sup> According to Kit Fine (1975) the vague predicate is the clearest case of vagueness. Characterizations of vagueness in other expressions (names and quantifiers) rest upon an account of meaning, but are arguably reducible to predicate vagueness. “For possibly one can replace, without any change in truth-value, each vague name by a corresponding vague predicate and each quantifier over a vague domain by an appropriately relativised quantifier over a more inclusive but precise domain” (p. 267). We can replace the name ‘Tibbles’ with a corresponding vague predicate, e.g. ‘... is Tibbles’, without change in truth-value. Likewise, a quantifier over the vague domain of all bald men can be replaced by a quantifier over everything that satisfies the vague predicates ‘... is bald’ and ‘... is man’ in a wider, but sharp, domain, e.g. of all physical objects (given that this domain is sharp). I agree with Fine’s account, and will consequently focus on vague predicates.



---

proposition to come out as either true or false.

It is worth noting that borderline cases are not only associated with vague predicates. For example some trolley cases in ethics, where you must decide between two actions with equally unfortunate outcomes, have borderline cases that do not depend on vagueness; one action is just as right or wrong as the other, making it impossible to determine what the 'right' choice is. Having borderline cases might be a necessary condition for vague predicates, but it is certainly not sufficient.

The problem with tolerance is interesting; it suggests that it is impossible to draw a precise boundary around the extensions of a vague predicate. This is related to the problem that Frege identifies in the opening quote: we cannot enumerate bald people, because we cannot identify the least bald people in a way that legitimately separates them from the least non-bald people (or the least indeterminately bald people). The problem is semantic; it is part of the meaning of 'bald' to require a significant difference between bald people and other things. No fact about the physical states of people's heads can help us make a sharp distinction.

Vague predicates, such as 'heap', 'bald' and 'red', have no sharp threshold for their application. Two people with approximately the same amount of hair on their heads are both either bald or non-bald, although one of them might have a few hairs less than the other. Whatever we consider the nature of the meaning of 'bald' to be, an overwhelming majority of speakers will agree that Patrick Stewart is bald, and that Ian McKellen is not bald. From this we can conclude that there is something bald, and something not bald. But the community of speakers cannot (I suppose) tell *exactly* where the line is between baldness and non-baldness. How much hair must Bob lose before he turns bald? Surely, not as much as Patrick Stewart, because Stewart could grow a couple of extra hairs and remain bald. Somewhere between McKellen and Stewart, but nowhere near either of them, people turn bald.

From the above account, it seems possible that Bob is in a state such that he is one split hair away from being bald. But here people, including many philosophers, will (and should) object. This is crazy! What we mean by 'bald' is not something that is determined by a single hair. If you *are* bald, growing

a single additional strand of hair will not change that. And if you are *not* bald, losing one single strand of hair will not change *that*. As long as the difference between two cases is small enough, the same predicate will apply equally to both cases. We call this the principle of tolerance.

**Tolerance:** Whenever  $F$  is vague and  $a$  and  $b$  are sufficiently similar with respect to  $F$ , then  $Fa$  is true if and only if  $Fb$  is true.

Once the principle of tolerance is observed, the following argument follows, known as a *sorites argument*:

(*Sor*) Someone with 200,000 hairs on his head is not bald.  
If someone with  $n$  hairs on his head is not bald, then  
    someone with  $n - 1$  hairs on his head is not bald.  
-----  
Someone with no hair on his head is not bald.

The argument is simple, and apparently valid. Still, most of us find the conclusion absurd because it contradicts our strong inclination to think that people like Patrick Stewart are clearly bald. This argument undermines the idea that we are using the predicate ‘bald’ to make true and meaningful descriptions of people. The obvious culprit is the principle of tolerance. Since the loss of one single hair is clearly an insignificant difference with respect to baldness, the principle of tolerance can account for every step from 200,000 to zero. But the loss of 10,000 hairs is not insignificant—regardless of whether the hairs were plucked one by one. Even so, it just does not seem right to deny the principle of tolerance, since that enables a situation in which the loss of one single hair makes Bob bald. And that is equally absurd.

Sorites arguments are usually unwelcome and are considered paradoxical. All we have to do in order to generate a contradiction in (*Sor*) is to include the uncontroversial premise “Someone with no hair on his head is bald.” This premise is inconsistent with the conclusion of the argument—and we have a paradox. The sorites come in many forms, and to the fascination and frustration of philosophers they all seem amazingly hard get rid of. Various responses have been proposed against the arguments, but still the sorites remain a lurking threat to our notions of truth and meaning of vague expressions.

---

A noteworthy criticism is directed against the epistemic character of the sorites arguments. In the borderline area, we have no independent justification for making a judgement; we rely solely on deductions from our previous judgements. This criticism is discussed in chapter 4. While it is likely the case that sorites-susceptibility is not a unique trait of semantic vagueness, it has alarming consequences. Due to sorites reasoning, vague predicates are in danger of exhausting their meaning. Sorites arguments go beyond the borderline cases; they question the cases that we believe to be clear.

The situation pulls in two opposite directions. We have strong motivation for resolving the vagueness of predicates in a way to avoid paradox, but also to dismiss vague predicates as defective and prevent them from ever entering the discourse. The latter strategy is briefly suggested by Ludwig Wittgenstein (1974). He claims that the problem with vague predicates is that they have no definite upper or lower limit. “I could decide: whatever is bigger than  $K$  cubic metres is to be a heap, and in that case the expression ‘The smallest heap’ has no meaning” (p. 240). The stipulation gives only one condition for being a heap; it says nothing about whatever is smaller than  $K$  cubic metres. Of course, as Wittgenstein also notes, the actual predicate ‘heap’ has an even less precise sense. The point is that it makes no sense to ask how small a heap can be, or how bald the least bald man must be; there is no lower limit.

“Make me a heap of sand here.” — “Fine, that is certainly something he would call a heap.” I was able to obey the command, so it was in order. But what about this command “Make me the smallest heap you would still call a heap”? I would say: that is nonsense; I can only determine a *de facto* upper and lower limit. (p. 240)

Since vague predicates lack *de facto* upper or lower limit, inquiries for such limits are misguided. Therefore, sorites arguments in themselves are nonsensical. They appear to consist of meaningful and well-formed sentences, but, like the making of the smallest heap, their meanings are impossible to determine in practice.

While this strategy effectively fights off the sorites paradox, we are left with the problems of tolerance and borderline cases; perhaps even to a greater extent

than before, since we have now committed to the non-existence of boundaries for vague predicates. Conceptually, there is no smallest heap, no greenest yellow, etc. An upshot with this kind of view is that it takes vagueness seriously; the predicate ‘heap’ is *really* vague, in every relevant sense of the word. Frege, who is a strong influence on Wittgenstein, argues that predicates without sharp boundaries must be disqualified from logical discourse.

A more popular strategy is to resolve vagueness by treating vague predicates as if they were sharp. The idea is that we can make predicates more precise through interpretations, or *precisifications*, that is a version of the vague predicate in which the borderline cases are determined as either applying or not applying. Thus, vague predicates are able to figure in propositions and logical inferences as if they were sharp predicates. While this makes the exact location of the borders more or less arbitrary, it enables us to express meaningful statements about what we believe to be true; for instance that Patrick Stewart is bald, that bananas are yellow, and that life has intrinsic value.

Under this strategy, the lack of a smallest heap is explained as some kind of ‘semantic laziness’; there are many possible precisifications of ‘heap’, but the community of speakers have not yet bothered to agree upon any single one of them as the correct meaning of ‘heap’. Speakers are therefore free to interpret ‘heap’ however they want, within the range of possible (and admissible) precisifications. David Lewis (1986) argues for the view that vagueness is such semantic indecision between precisifications: “The reason it’s vague where the outback begins is not that there’s this thing, the outback, with imprecise borders; rather there are many things, with different borders, and nobody has been fool enough to try to enforce a choice of one of them as the official referent of the word ‘outback’.” (p. 212).

Borderline cases arise when the community of speakers disagree—where the precisifications diverge. Thus, borderline cases (and vagueness) exist between precisifications, and not within them. Tolerance is weakened to a global principle that simply follows from the fact that there are borderline cases; if there is an insignificant difference in Bob and Al’s baldness, then it cannot be the case that all speakers agree that one of them is bald and the other is non-bald. This seems to block the sorites paradox. (We consider how in chapter 1.)

---

The aim of this thesis is to assess the foundations of the popular strategy, particularly the notion of precisifications. I have two worries. The first is that while precisifications may prove valuable for simulating vagueness in formal languages, they fail to capture the true nature of vagueness. I will argue that this strategy, if ever accomplished, will not explain vagueness, but simply dissolve it. My second worry is that the notion of precisifications is too unclear. While there are different ways of cashing out the term, they all seem to rely on some further vague notion. I will show that a precise notion of precisifications depends on a strict set-theoretical understanding of the extensions of predicates, which in turn lacks the robustness and flexibility that seem to characterize predicates.

The first worry is baseless to those who believe that vagueness is nothing more than semantic indecision, but the second worry is crucial. If it should turn out that a precise notion of precisifications is unwarrantable, then there might be more than semantic indecision to vagueness after all.

## PROCRUSTES

According to Greek legends, the man known as Procrustes, ‘the stretcher’, lived on a mountain on the way between Athens and Eleusis. He had the habit of inviting travellers to spend the night in his iron bed. Procrustes wanted the bed to fit his visitors to perfection, and he made sure that it did in the most cruel way by stretching his guests out with a smith’s hammer, or, if the guest proved too large for the bed, by amputating any excess limbs.<sup>3</sup> The phrase ‘Procrustes bed’ is later used of situations in which things are required to fit perfectly within some arbitrary standard.

Besides the mutilation and killing of innocent travellers, the procrustean approach strikes me as fundamentally misguided. Although he does provide his guests with a bed that fit them perfectly, Procrustes turns the order of things upside down by altering the “wrong” variable. Apparently it is categorically wrong to physically change people to make them fit into beds, hats, shoes, and so on.

---

<sup>3</sup>In the more sinister tales, Procrustes has two beds—one large (for short guests), and one small (for tall guests)—to ensure that nobody would fit the bed. See Karl Kerényi (1959, pp. 222–223).

In the people-fitting-to-beds relation, people are more fundamental. If there is no fit, it is the bed that has the wrong shape, not the person. A similar principle exists in philosophy of science: it is considered bad scientific practice to alter data to fit your theory, for instance by rejecting parts of some statistical data because it contradicts your theory. Ideally, the theory should fit the data, not the other way around.<sup>4</sup> (Let us refrain from discussing this last claim. There are interesting and important theoretical and practical issues with this view, but unfortunately that is way out of our current scope.) The general idea is that the relation ‘*a* fits *b*’ is in practice asymmetrical—although, strictly speaking, it is not.

In the case of vagueness, Procrustes is the theorist who wants vague predicates to fit the aspects of his existing theory of predicates. This does not mean that there must be one true bed for every predicate; language-Procrustes might have several beds for various purposes and contexts. The point is that the predicate must be made to fit a bed perfectly before it makes sense to him. This is quite understandable, since vague predicates have no determinate shape in themselves. Even in his least homicidal moments, Procrustes will not be able to provide a bed that would fit a vague predicate. Not without some light hammering.

Procrustean approaches to vagueness is widespread, probably due to the intimacy of sharp predicates and classical logic. Our default assumption is that predicates are sharp. Some well-known slogans include: ‘*p* or not *p*’ (“You are either too sick for school, or you are not.”), and ‘if every *F* is *G* and *a* is not *G*, then *a* is not *F*’ (“If every film by Woody Allen is worth seeing and ‘Glitter’ is not worth seeing, then ‘Glitter’ is not a film by Woody Allen.”). These are valid schemes, but they presuppose that predicates are sharp. If the predicate ‘worth seeing’ is vague, as it probably is, then we cannot draw a line between ‘Manhattan’ and ‘Glitter’ in respect to their worthiness of seeing without violating Tolerance. A procrustean solution to this problem is to tie the predicate ‘worth seeing’ to a Procrustes bed; any size will do, because it will enable us to make

---

<sup>4</sup> It is also considered bad practice to alter your theory to fit your data, but that is more of a methodological heuristic than a fundamental mistake. A theory is more likely to be false if it only predicts what you have verified in advance. Nevertheless, a theory *is* false when it is falsified by data—whether you choose to ignore that data or not.

---

sense of the predicate. In order to be useful, the predicate ‘worth seeing’ must be true for some things and false for others.

Whether a procrustean approach provides a good analysis of vagueness will be an underlying question throughout this thesis. There is no doubt that vagueness is easier to handle in the hands of Procrustes; his solution is both plain and simple. But despite the success rate, we have good reason to doubt this as the right approach to vagueness—particularly since it involves handing over parts of our language to a madman.





## Chapter One

# *The Sorites Paradox*

*The sorites paradox is considered a paradigmatic case of vagueness, and it is perhaps one of the most vicious problems in philosophy of language. ‘Sorites’ is derived from an Ancient Greek word for ‘heap’, which refers to the first known sorites: the problem of how many grains of sand it takes to make a heap. In this chapter I present different forms of the sorites paradox, and some notable responses to them.*

### 1.1 VARIOUS FORMS OF THE SORITES PARADOX

The most simple form of the sorites paradox is the conditional. It consists of a series of conditional statements, which leads from an apparently true premise to an apparently false conclusion. The paradox arises when there are vague predicates, like ‘bald’, in our premises. We hold it to be true that Patrick Stewart is bald, and false that Ian McKellen is bald. These are both what we may call clear cases of baldness; there is no doubt as to whether the predicate applies. Patrick Stewart is clearly bald, and Ian McKellen is clearly not bald. Yet, through a series of small, insignificant variations, we are able to produce the following argument, which contradicts our initial beliefs.

Let  $F$  be the predicate ‘bald’, and  $\langle a_1, \dots, a_n \rangle$  be an ordered series of people (where  $n$  is an arbitrary number)—starting with Stewart ( $a_1$ ), and ending with McKellen ( $a_n$ ). The people in the series are ordered such that  $a_1$  differs from  $a_2$  by one single hair,  $a_2$  differs from  $a_3$  by another single hair, and so on. Since the predicate  $F$  is *tolerant* such that variations of one single hair are insignificant with respect to whether the predicate  $F$  applies, we can form a long series of

## 1. THE SORITES PARADOX

---

conditional statements:<sup>1</sup> if  $a_1$  is bald then  $a_2$  is bald, if  $a_2$  is bald then  $a_3$  is bald, and so on. From the premise that Patrick Stewart is bald, and the series of conditional statements, it follows that Ian McKellen is bald as well, which we believe he is not. Formally, we can set up the conditional sorites argument this way:

$$\begin{array}{l} (\text{SorCond}) \quad Fa_1 \\ \quad \quad \quad Fa_1 \rightarrow Fa_2 \\ \quad \quad \quad Fa_2 \rightarrow Fa_3 \\ \quad \quad \quad Fa_3 \rightarrow Fa_4 \\ \quad \quad \quad \dots \\ \quad \quad \quad \frac{Fa_{n-1} \rightarrow Fa_n}{Fa_n} \end{array}$$

Such an argument is ‘soritical’ if the series  $\langle a_1, \dots, a_n \rangle$  is ordered with respect to  $F$ , and each member of the series is sufficiently similar to its neighbours for the tolerance of  $F$  to apply.<sup>2</sup> Further,  $Fa_1$  must appear to be true, and  $a_n$  must be sufficiently different from  $a_1$  (with respect to  $F$ ) so that  $Fa_n$  appears to be false.

This form of the argument, known since antiquity, is also called a little-by-little-argument. It is quite clear that the subject is lead into accepting a conclusion she believes to be false, by adding up a large amount of small, insignificant changes with respect to  $F$ . Yet, since we accept the tolerance of  $F$ , and every pair of adjacent  $a$ -s are sufficiently similar for tolerance to apply, there seems to be no way of calling the bluff.

A different form of the paradox, the mathematical induction sorites, involves a generalized version of the principle of tolerance as its inductive step.<sup>3</sup> The argument shows that if  $F$  is true of some  $a$ , then it is true of every  $a$ .

---

<sup>1</sup>The conditional series is in principle infinite, but it is sufficient for the argument that it is construed for the first  $n$  instances.

<sup>2</sup>This constraint serves to validate all the conditional premises in the argument, as noted by Jonathan Barnes (1982). If  $F$  is the predicate ‘bald’, then every pair of adjacent  $a$ -s might differ with one hair, if that is sufficiently small for  $F$  to apply to both  $a$ -s in every case, i.e. one single hair cannot determine whether you are bald.

<sup>3</sup>The principle of tolerance is presented in the introduction, and is further discussed in the upcoming section. The principle is, basically, that vague predicates are insensitive to tiny differences, such that tiny (insignificant) differences cannot determine whether a vague predicate applies or not.

$$\begin{array}{l}
 (\text{SorMath}) \quad Fa_1 \\
 \quad \frac{\forall i(Fa_i \rightarrow Fa_{i+1})}{\forall i(Fa_i)}
 \end{array}$$

This argument is soritical and valid under the same constraints as the conditional form.

Last we have the line-drawing sorites, which directly contradicts the principle of tolerance. We start from the premise that  $Fa_1$  is true, and the denial of the apparently absurd conclusion from (*SorMath*): that every  $a$  is  $F$ . From this the conclusion follows that there must be an adjacent pair of  $a$ -s where  $F$  is true of one and false of the other; in other words that tolerance is false.

$$\begin{array}{l}
 (\text{SorLine}) \quad Fa_1 \\
 \quad \frac{\neg \forall i(Fa_i)}{\exists i \geq 1(Fa_i \wedge \neg Fa_{i+1})}
 \end{array}$$

The conclusion is that there must be a sharp cut-off in the series  $\langle a_1, \dots, a_n \rangle$ , if  $F$  is such that it is true of something, but not true of everything. The reasoning seems quite sound, and the conclusion seems acceptable as a general rule. Yet, with respect to vague predicates, it seems absurd. As Graham Priest (2003) describes the sorites phenomenon: “It arises simply because we are forced to recognize the existence of cut-off points where both common sense and philosophical intuition scream that there are none” (Priest 2003, p. 10).

We may ask whether the sorites arguments really are paradoxical. There are at least three possible answers.<sup>4</sup> The sorites may be *falsidical paradoxes*, i.e. false arguments that only look sound to us because we have yet to unveil their underlying fallacy or false premise. An adequate response to such a paradox must explain what kind of fallacy we commit, or which premise is false, and why. The multitude of different formulations of sorites arguments makes such a response extremely difficult, if not impossible. Another strategy is to accept the sorites as *veridical paradoxes*, i.e. true arguments, that only appear to be absurd. Such a response must make an effort to explain why we find the arguments counter-intuitive. This quickly leads down the paths to the extreme responses; either

---

<sup>4</sup>Following the taxonomy of paradoxes suggested in W. V. Quine (1962), which is also an excellent introduction to paradoxes in general.

## 1. THE SORITES PARADOX

---

denying the phenomenon of (semantic) vagueness, or denying that truth applies to vague expressions at all. The final, and perhaps most frequently pursued option is to accept that the sorites are real paradoxes, also known as *antinomies*, which simply cannot be resolved within classical logic and the current leading theories of truth and meaning.

Although these paradoxes at first glance might seem both frivolous and fun puzzles or mind-games, the paradoxes of vagueness are considered to be some of the most vicious philosophical problems. These problems have received fluctuating attention throughout history, but they keep coming back to haunt us. As language speakers we are daily faced with problems of vagueness, and yet we are somehow able to ignore most of these problems. Apart from this blissful ignorance in our everyday life, there is little agreement between the different philosophical approaches to vagueness.

### 1.2 STRONG AND WEAK TOLERANCE

This is the strong formulation of the principle of tolerance:

(*Tol*) Whenever  $F$  is vague and  $a$  and  $b$  are sufficiently similar with respect to  $F$ , then  $Fa$  is true if and only if  $Fb$  is true.

Tolerance is a semantic principle governing some specific kind of predicates: the vague (or tolerant) predicates. We call it ‘semantic’ because it directly affects the truth-conditions of  $F$ . If  $F$  is tolerant, then a substantial difference between two instances is required for one of them to be  $F$  and the other to be not  $F$ . For example whenever something is red, a slight variation in tint is not sufficient to say that the resulting colour is no longer red. Yet, a large change in tint is sufficient to say that the resulting colour is no longer red, but e.g. orange. The predicate ‘red’ is such that the two tints of red, despite their slight (insignificant) difference, are both red in the same sense of ‘red’. A response to the sorites which suggests that every different shade of red has its own predicate, such that  $a$  is reflecting-light-of-650-nm-wavelength-kind-of-red and  $b$  is reflecting-light-of-651-nm-wavelength-kind-of-red, is hard to justify, at least intuitively, because it implies that our use of language makes distinctions and precise measurement

that we—as human speakers—cannot be expected to carry out with our perceptual abilities and cognitive capacity.

Crispin Wright (1975) explains the tolerance of vague predicates as a feature essential to their explanatory role. We attribute different properties and moral status to children versus adults. In order for this distinction to be useful (and morally defensible), there must be a substantial difference in maturity between children and adults.<sup>5</sup> Therefore, the predicate ‘child’ is of no use to us if it is sensitive to tiny, insignificant changes in maturity. Children do not lose their moral status as children from one heartbeat to the next.

A lot of predication in ordinary discourse is based on casual observation. We see whether a car is red, we hear whether a sound is loud, and we taste whether a soup is salty. Therefore, useful predicates cannot make distinctions that members of the community of competent speakers cannot observe, i.e. two indistinguishable patches of colour, or two bowls of soup with an undetectable difference in saltiness. Too fine-grained predicates in these respects, if they ever existed, would have been washed out of our language a long time ago; since they were impossible to learn, apply or understand.

Although the case for a tolerance principle is strong, some consider the formulation (*Tol*) to be too strong. There is a problem that (*Tol*) generates a paradox. It secures the inductive step in soritical reasoning, like in this argument from last section:

$$\begin{array}{l}
 (\text{SorMath}) \quad Fa_1 \\
 \frac{\forall i(Fa_i \rightarrow Fa_{i+1})}{\forall i(Fa_i)}
 \end{array}$$

The second line in the argument is the crucial one. According to Diana Raffman (1994) the strong principle of tolerance ignores an important feature of the sorites series: Every item in the ordered series is sufficiently *F*-wise similar to *the next* item for *F* to apply the same way to both of them. Non-adjacent items (in the ordered series) may be significantly different so that *F* does not apply the same way to both of them. The inductive premise, and hence also the principle of tol-

<sup>5</sup>I do not speak of the legal definitions of children and adults—which are more or less arbitrary conventions.

## 1. THE SORITES PARADOX

---

erance, is acceptable only if adjacent members in the series are judged pairwise. This suggests a weaker principle of tolerance:

(*TolC*–) Whenever  $F$  is vague and  $a$  and  $b$  are sufficiently similar with respect to  $F$ , then  $Fa$  and  $Fb$ , when judged pairwise, will be assigned the same truth-value.

The distinctive feature of pairwise judgement is that the lack of relevant difference between  $a$  and  $b$  is clearly presented to the subject; there is no justification in the sample for making the cut right there. If  $a$  looks red and  $b$  looks just like  $a$ , then  $b$  looks red too, provided that  $a$  and  $b$  are judged pairwise. This relation of similarity holds between pairs of adjacent cases, but is not transitive. Although  $a_1$  is similar to  $a_2$  and  $a_2$  is similar to  $a_3$ , it does not follow that  $a_1$  is similar to  $a_3$ , as illustrated by the diagram in Figure 1.1.



Figure 1.1: Diagram illustrating the non-transitivity of a similarity relation.

If we suppose that  $a_1$  is clearly red and  $a_n$  is clearly orange, we may accept that every pair of adjacent items are sufficiently similar for tolerance to apply. But this does not mean that every pair (adjacent or not) is sufficiently similar for tolerance to apply. E.g.  $a_n$  is clearly orange and thus significantly different from the clearly red  $a_1$ . When judged individually, the similarity relations play no role, and the cases are judged independently of their position in the sorites series.

According to Stewart Shapiro (2008), vague predicates have a feature he calls ‘open texture’: their borderline cases can be decided by a speaker in context. If a thirteen-year-old is a borderline child, the speaker is free to classify her as a child, or as not a child (or not classify her as a child/non-child at all).<sup>6</sup> This classification is contextual, and any shift in context releases the speaker from

---

<sup>6</sup>There are some conceptual restrictions on the speaker, sometimes called the *penumbral connections*. For instance if the speaker decides that Harry is bald, then (in that context) everyone as bald as or balder than Harry must be bald too. Penumbral connections are further discussed in section 1.5.

her previous commitments. It is important to notice that it is not the meaning of vague predicates that varies. The open texture is part of the meaning of vague predicates. What varies in context is what the speaker decides to count under the extension of the vague predicates. In other words: the extensions varies with context.

The contextualist approach to the sorites paradox is thus to deny the inductive step. It is not true that the judgement that  $a_1$  is red commits the speaker to the redness of  $a_n$ . Somewhere in the borderline region the speaker will be presented with a pair of two almost clearly orange items, e.g.  $\langle a_{249}, a_{250} \rangle$ , and decide that they are not red. Although she has already said that the pair  $\langle a_{248}, a_{249} \rangle$  looks red, the context has just shifted such that this previous verdict is overruled by her most recent judgement. Consequently,  $a_{249}$  (and possibly  $a_{248}$  and some other items as well) are no longer counted under the extension of the predicate 'red'. This strategy is compatible with the weak principle of tolerance, as formulated above.

### 1.3 MARGIN OF ERROR

Another response is to interpret the sorites as arguments against the principle of tolerance. The arguments prove, contrary to our expectations, that there is a sharp cut-off in the application of a vague predicate. The insignificant differences in a sorites series are only apparently insignificant. There *is* in fact one pair of adjacent  $a$ -s such that  $F$  is true of one and false of the other, even though we cannot know which pair this is. (And even if we did, we would not have been able to observe the tiny, yet substantial difference.) The strong appearance of tolerance is a result of our necessarily limited knowledge of vague concepts. The actual concepts behind words like 'bald', 'heap', 'many', 'child', and so on, are too complex for human beings. Vagueness is a matter of ignorance. Such a view on vagueness is defended by Timothy Williamson:

You have no way of making your use of a concept on a particular occasion perfectly sensitive to your overall pattern of use, for you have no way of surveying that pattern in all its details. Since the

## 1. THE SORITES PARADOX

---

content of the concept depends on the overall pattern, you have no way of making your use of a concept on a particular occasion perfectly sensitive to its content. (Williamson 1994, pp. 231–232)

Because human beings are cognitively disabled from using so-called vague concepts in a perfectly precise manner, our use is always subject to a certain margin of error. What I mean by ‘bald’ when I say that Bob is bald, cannot be that Bob is bald to the precision of one single hair. This is because I am unable to provide that level of precision. I am neither able to perceive nor picture Bob as having exactly  $n$  number of hairs. Consequently, if I say ‘Bob is bald’, then him growing one single hair will not change the truth of my proposition.

The margin of error principle is different from the principle of tolerance. Tolerance is a semantic principle concerning the meaning of expressions, whereas margin of error is an epistemic principle concerning what we know and what kind of judgements we can make.

- (!) Whenever  $F$  is vague,  $a$  and  $b$  are sufficiently similar with respect to  $F$ , and we know that  $Fa$  is true, then  $Fb$  is true.

This principle says that if  $Fa$  is known—if it is both true and you know that it is true—then we can infer by the margin of error that the adjacent  $Fb$  in the sorites series is true as well. This does not say, as the principle of tolerance does, that the truth of  $Fa$  entails the truth of  $Fb$ . What it says is that the *known* truth of  $Fa$  entails the truth of  $Fb$ . Since we are unable to use  $F$  with the precision required to discriminate  $Fa$  and  $Fb$ , we cannot know that  $Fa$  is true unless it is safely within our margin of error, in which case  $Fb$  is true as well.

We might worry that (!) gives rise to an epistemic version of the sorites paradox: If it is known to be true that Bob is bald when he has  $n$  hairs, then it is true that Bob is bald when he has  $n + 1$  hairs, therefore it is known to be true that Bob is bald when he has  $n + 1$  hairs, and it will also be true that Bob is bald when he has  $n + 2$  hairs, and so on. Ultimately, since everyone has a number of hairs, if it is known that someone is bald, then it must be known that everyone is bald. But this line of argument needs one further assumption: the *KK-thesis*, which is false.



(*KK*) If you know  $p$ , then you know that you know  $p$ .

Without (*KK*) it does not follow from (!) that it is known that  $Fb$ , because the fact that it is known that  $Fa$  only entails that  $Fb$ , and not that it is also *known* that  $Fb$ . Hence the paradox is blocked.

Williamson (2000) provides the following argument against the *KK*-thesis. Consider a man, Mr. Magoo, looking at a tree. By looking at the tree, he cannot know exactly how tall it is. Even if he judges that the tree is 550 cm tall—and that is indeed the height of the tree—he does not then *know* the height of the tree; he has only guessed it correctly. For all Mr. Magoo knows, the tree might be 551 or 549 cm tall. He knows this, hence for any natural number  $n$

( $I_n$ ) Mr. Magoo knows that if the tree is  $n + 1$  cm tall, then he does not know that the tree is not  $n$  cm tall.

We assume that knowledge is closed such that if Mr. Magoo knows the set  $X$  of relevant propositions, and  $X$  entails a proposition  $p$ , then Mr. Magoo knows  $p$  (after some reflection over  $X$ ).<sup>7</sup>

( $K_1$ ) Mr. Magoo knows that the tree is not 1 cm tall.

It is clear to Mr. Magoo that the tree must be way more than 1 cm tall. From ( $K_1$ ) and (*KK*) we can infer that

( $KK_1$ ) Mr. Magoo knows that he knows that the tree is not 1 cm tall.

By instantiating ( $I_n$ ) we can say that

( $I_1$ ) Mr. Magoo knows that if the tree is 2 cm tall, then he does not know that the tree is not 1 cm tall.

If the tree was 2 cm tall, it would follow that Mr. Magoo did not know that the tree is not 1 cm. But this can be ruled out by ( $K_1$ ). And since Mr. Magoo is aware of this, from ( $KK_1$ ) and ( $I_1$ ) it follows that

---

<sup>7</sup>This is not a claim that we know all the consequences of the things we know (that is absurd), but that reflective deduction is possible within the domain of knowledge. We assume that Mr. Magoo in the argument draws the relevant conclusions about his own state.

## 1. THE SORITES PARADOX

---

( $K_2$ ) Mr. Magoo knows that the tree is not 2 cm tall.

This can be repeated until it reaches the actual height of the tree. But Mr. Magoo cannot know that the tree is not 550 cm, since that would be false and knowledge is factive—at least according to Williamson. This means that there is a false assumption in the above argument; Williamson points to the KK-thesis. We do not always know that we know things. This is not because we are reflectively unaware of it or because the closure of knowledge is false, but because sometimes our knowledge is inexact. “To the informed observer, hearing gives some knowledge about loudness in decibels, and touch about heat in degrees centigrade” (Williamson 2000, p. 119). Mr. Magoo’s observations of the tree gives him some knowledge of its height in centimetres, but there is no least number  $n$  such that he is aware that he knows that the tree is taller than  $n$  cm, but not that the tree is taller than  $n + 1$  cm.<sup>8</sup> We cannot identify exactly where (KK) fails,<sup>9</sup> but, according to Williamson, by generalizing this argument we see that it does.

One upshot of the epistemic view on vagueness is that we are able to resolve the sorites paradoxes by tossing out the cumbersome tolerance principle. We may preserve classical logic and continue working within the standard theories of truth and meaning of language, in which the exact truth-conditions of every meaningful expression can (in principle) be given. Vagueness is a non-paradoxical epistemic phenomenon that is safely contained within our representational world, there is no real threat to the foundations of language and concepts—nor to the external (real) world.

An influential criticism against epistemicism, raised by Wright (1995), is that the epistemic view provides no clues as to what kind of reference that is supposed to hold between vague expressions and the properties or objects they pick out. Since the epistemicists holds that e.g. an utterance of ‘TW is thin’ is true or false of TW in the (external) world, then ‘thin’ must pick out some property held by TW and anyone else with equal or less than some exact physical measurements  $m$ . Thus ‘thin’ must refer to the property of having physical measurements equal to or less than  $m$ . If TW is borderline thin, then speakers

---

<sup>8</sup>This is, very roughly, Williamson’s (1994) explanation of vagueness.

<sup>9</sup>As Williamson (2000) notes, I cannot give a counterexample to (KK), because that would require me to know of a  $p$  that I know, of which I do not know that I know.

cannot know whether TW has this property; not because they do not know the relevant physical measurements of TW, but because they do not know what kind of property ‘thin’ refers to. On the other hand, if TW was clearly thin, speakers would still not know which property ‘thin’ refers to (but somehow they would still know that TW was in possession of it). There is an air of what Wright calls “semantic mysticism” surrounding the epistemic view on vagueness. The epistemicist cannot possibly, by our use of ‘thin’, account for the meaning of that word. At least not in terms that would suggest that the word has a precise meaning.

The above picture is further complicated by the fact that speakers may introduce vagueness intentionally. If I ask whether it is a long train ride from Oslo to Stockholm, I do not mean ‘long’ in any precise (i.e. sharp) sense. Since most people will be able to understand how to respond appropriately to my question, they must somehow grasp what I mean by ‘long’, viz. the intentionally vague sense. Likewise when I say ‘I’ll be back in about an hour’, I do not mean by ‘about an hour’ a precisely delimited period between, say, 50 and 70 minutes from now. In these cases lack of knowledge and margin of error are unfit to explain the semantic tolerance of the vague expressions. The vagueness is part of the respective meanings of the expressions, and not merely an epistemic safeguard against believing something that we are unable to determine whether is really true or false.

#### 1.4 NIHILISM

I cannot deny that tolerance has a strong intuitive appeal. From the conditional sorites in section 1.1, we have seen that, in a sorites series there is no pair of adjacent cases that invalidates the tolerance principle. But combined with the assumption that something is  $F$  and something is not  $F$ , tolerance leads to paradox. This makes us question our inclination to accept the tolerance principle. One response strategy, as we have seen, is to construct weak alternatives to the principle of tolerance that block the sorites paradox, but preserve parts of our inclination towards tolerance. But there are radical options to this approach. The epistemic view, that we discussed in the last section, rejects tolerance entirely as

## 1. THE SORITES PARADOX

---

unsound. Another radical approach is to read the sorites not as a logical paradox, but rather as a direct or indirect argument against the view that vague predicates reflect real distinctions in the world. Tolerance makes vague predicates incoherent; if  $F$  is sorites-susceptible, then  $F$  is incoherent (as we can demonstrate by a sorites argument). This is called ‘semantic nihilism’.

Peter Unger (1979) claims that we can accept the conclusion of sorites arguments without facing any important logical problems. Let us consider one of his arguments for nihilism. It begins with our notion of *ordinary things*, those things that we experience in our everyday life: tables, chairs, stones, twigs, etc. From among those things, we choose an arbitrary object: a stone.

- (1) There is at least one stone.
- (2) For anything there may be, if it is a stone, then it consists of many atoms but a finite number.
- (3) For anything there may be, if it is a stone (which consists of many atoms but a finite number), then the net removal of one atom, or only a few, in a way that is most innocuous and favorable, will not mean the difference as to whether there is a stone in the situation.<sup>10</sup>

These premises are inconsistent; because of (1), (3) is in conflict with (2). There cannot be a stone that must consist of many atoms, but can always survive the removal of one atom, or only a few. Sooner or later there cannot be enough atoms to make a stone according to (2), although (3) would still insist that there is a stone. Since both (2) and (3) seems necessarily true given our concept STONE, Unger concludes that (1) must be the least plausible premise. Therefore

- (4) There are no stones.

Since the stone is arbitrary, the argument goes equally for any other ordinary thing and we may generalize (4)

- (5) There are no ordinary things.

---

<sup>10</sup>Unger 1979, p. 120.

There is a clear soritical structure to this argument.<sup>11</sup> (3) expresses the tolerance principle for the concept STONE; stones are essentially such that they can survive the net removal of one single or a few atoms—anything else seems absurd. Yet the conclusion is rather counter-intuitive as well.

It is important to notice that the argument does not claim that if I hold a stone in my hand, what I am holding in my hand is something that does not exist, or that I am in fact not holding anything in my hand. Rather, the conclusion is that whatever is in my hand, it does not fit the concept STONE, nor the concept of any other ordinary object. The idea is that nothing can fit such concepts, because the concepts themselves are incoherent (and we assume that things in the world are not incoherent). Accepting Unger's argument does not in itself commit us to a completely different outlook on the world (although some might insist on talking about 'collection arranged in a stone-like manner' where the ordinary term 'stone' would suffice). What we give up is the assumption that all predicates in a language can be mapped to properties of objects and relations between objects. That I am holding a stone in my hand is only true or meaningful in an approximate sense; to the extent that whatever is in my hand might appear to be stone-ish.<sup>12</sup> Classical logic is not fit to recognize approximate truth and meaning, thus we cannot talk about whether or not 'This is a stone' is true.<sup>13</sup> As far as logic is concerned, there are no stones. Hence we can (and should) accept the sorites arguments as sound without facing any important logical problems.

In order to save ordinary concepts from the sorites we must, according to Unger, believe in one of the two following "miracles". The first miracle is a metaphysical one, *the miracle of metaphysical illusion*: There is literally a limit to how many atoms that can be removed from a stone. When that limit is reached, the removal of more atoms either becomes physically impossible or causes a drastic change in the remaining object; it ceases to exist or turns into something completely different. But this is incompatible with both our everyday

---

<sup>11</sup>Unger calls this argument 'the sorites of decomposition'.

<sup>12</sup>By 'extent' I do not mean to apply to a fuzzy or continuum-valued logic. I mean something closer to the informal phrase 'more or less'. The claim might, however, be translated to a (more or less) adequate expression in such a logic, as opposed to classical logic which would require a strict boolean truth-value (i.e. *true* or *false*).

<sup>13</sup>See David Braun and Theodore Sider 2007, for an argument for *approximate truth* as the correct notion of truth for vague expressions.

experience and our current scientific theories. Apart from our fondness of ordinary concepts we seem to have no reason to believe in it. The second kind of miracle is epistemic, *the miracle of conceptual comprehension*: The concepts we apply are indeed sharply delimited, and contrary to what most speakers seem to believe STONE is not indifferent to the net removal of one, or a few, atoms. This last miracle is what the epistemic view on vagueness proposes. Unger rejects both miracles as implausible, along with his belief in stones.

### 1.5 SUPERVALUATIONISM

Despite some unruly sorites examples, most of the time we seem perfectly capable of using and understanding vague predicates successfully in ordinary discourse. One explanation for this is that we are able to make vague expressions precise. I am able to assess whether Harry is bald, not because I am somehow able to think in terms of the incoherent predicate ‘bald’, but because I am able to derive a sharp predicate that resembles ‘bald’, except that it is determined for every case (and hence not incoherent). This is the idea of precisifications. According to Fine (1975) every vague predicate has a range of admissible precisifications, i.e. more precise predicates that does not conflict with the conceptual constraints of the vague predicate.<sup>14</sup> If Ian McKellen is not bald on any admissible precisification of ‘bald’, then he is definitely not bald.

Let  $F$  be a vague predicate,  $v_1, \dots, v_4$  are admissible precisifications of  $F$ , and  $\langle a_1, \dots, a_5 \rangle$  is a relevant sorites series. This gives the following table:

	$s$	$v_1$	$v_2$	$v_3$	$v_4$
$Fa_1$	⊤	⊤	⊤	⊤	⊤
$Fa_2$	–	⊤	⊤	⊤	⊥
$Fa_3$	–	⊤	⊤	⊥	⊥
$Fa_4$	–	⊤	⊥	⊥	⊥
$Fa_5$	⊥	⊥	⊥	⊥	⊥

---

<sup>14</sup> One example of such a conceptual constraint for e.g. ‘tall’ is that if Tally is tall, then everyone as tall as or taller than Tally are tall. A precisification that counts Tally as tall, but Haley (who is taller than Tally) as not tall is not an admissible precisification of ‘tall’. Fine calls these constraints *penumbral connections*.

Notice that  $Fa_1$  is true under any precisification of  $F$ , and  $Fa_5$  is false under any precisification of  $F$ . These are the definite cases, the others we may classify as borderline cases of  $F$ . The truth-values in the column  $s$  are only the definite cases, also known as the *supervaluations* of  $F$ . It is also worth noting that the precisifications respect the conceptual constraints of  $F$  such that they only make one cut in the series; if  $Fa_1$  comes out true and  $Fa_4$  comes out true, then  $F$  is true of all the other instances between  $a_1$  and  $a_4$  in the series as well.<sup>15</sup>

One of the strengths of supervaluationism is that logical sentences such as instances of the law of excluded middle comes out as true under any precisification. Although  $Fa_3$  does not have a truth-value under the supervaluation,  $Fa_3 \vee \neg Fa_3$  is supertrue (it is true under any precisification of  $F$ ).

At first glance the supervaluationist response to vagueness looks rather satisfactory. A case that comes out true under some precisifications of  $F$ , and false under others, are regarded as a borderline case; hence it is neither true nor false. We may consider borderline cases as representing a third kind of truth-value, the indefinite, or the truth-value “gap”. This requires us to abandon the principle of bivalence for supervaluations. It is not such that every statement is either true or false, and hence the non-truth of  $p$  does not entail that  $p$  is false (because  $p$  might be indefinite). Likewise the non-falsity of  $p$  does not entail  $p$ . When our semantics are appropriately tuned, we may weaken the tolerance principle accordingly.

(*TolS*–) Whenever  $F$  is vague and  $a$  and  $b$  are sufficiently similar with respect to  $F$ , then it is never such, under supervaluation, that  $Fa$  is true and  $Fb$  is false.

There is no case in which  $Fa$  is true and one of its adjacent instances (in a sorites series) is false. This means that if Harry is definitely bald, then Barry (who has only one single hair more than Harry) is either bald or borderline bald. This principle does not support the strong inductive reasoning that is required in order to make a sorites paradox. The truth of  $Fa$  does not entail the truth of its adjacent

<sup>15</sup>All vague predicates do not need to have this constraint on a sorites series. For instance the predicate ‘is a mountain’ might not present a clearly ordered series of instances with respect to their degree of mountain-ness. But this matter is not important to the current discussion, and would only serve to make the present picture more confusing.

cases, as the original (*Tol*) proposes. The weakened (*TolS*–) says only that the adjacent cases cannot be false—they may be true or indefinite. This appears to dissolve the sorites arguments, but it does not.

Consider the predicate ‘definitely bald’. We would expect this to be a sharp predicate, clearly dividing those that are definitely bald from everyone else (i.e. both the non-bald and the borderline bald). But apparently that is not the case; the situation in which the loss of a single hair determines whether or not you are definitely bald remains absurd and unacceptable. If the supervaluationist maintains that ‘definitely bald’ contains borderline cases, she is caught in a regress of higher-order vagueness. If she does not, she owes an explanation on why we find the above story so counter-intuitive.

### 1.6 OPPOSING VIEWS, YOU’RE BOTH RIGHT

Jc Beall (2010) proposes another way of solving the sorites paradox. The question of semantic nihilism divides philosophers into two camps: the semantic nihilists and the non-nihilists. Beall claims to provide a unified account of vagueness that preserves the best from both camps, without anyone having to give up their core beliefs. The nihilists are granted that vague predicates are radically tolerant and do not yield nice, sharp distinctions in the world; strictly speaking, they are useless. On the other hand, the non-nihilists are granted that vague predicates somehow describe the world; there is a sense in which vague predicates are useful after all. But the non-nihilists “may be seen as giving an account of the sharp offspring [i.e. precisifications] of our vague predicates, not an account of vagueness itself” (Beall 2010, p. 193). This suggests that the controversy concerning the sorites paradox is mostly due to a confusion of vague predicates and their precisifications.

A unified account dissolves the paradox by adding different flavours to the premisses in the sorites paradox; i.e. ‘5,000 grains of sand do, *in some sense*, make a heap’ (non-nihilist), and ‘ $n$  grains of sand do, *strictly speaking*, make a heap if and only if  $n + 1$  grains of sand do, *strictly speaking*, make a heap’ (nihilist). The fact that some number of grains do, in some sense, make a heap does not suffice to establish that the same number of grains do, strictly speaking,



make a heap, nor vice versa. Thus, we seem to avoid both the logical paradox, and the clash with intuition and common sense.

Beall's analysis is that the tension between nihilist and non-nihilist treatment of vagueness is caused by each of them strongly holding one of the following—apparently mutually exclusive—positions:

(*NoCut*) Vague predicates have no cut-offs. If one thing is counted into the (positive or negative) extension of a vague predicate, everything must be counted in, and vice versa.

(*Utility*) Vague predicates can be used to make true descriptions of the world.

The position (*NoCut*) follows from the nihilist conclusion of the sorites arguments: if one man is bald, then someone with one single strand of hair more than him is bald too, then, by multiple induction steps, everyone with more hair than the bald man are bald (and we already knew that everyone with less hair than the bald man are bald); hence, if someone is bald, everyone are bald. This proves that, since vague predicates either counts all in or all out, they cannot be used to make true descriptions of the world—of which we assume that it contains both bald and non-bald people. Hence (*Utility*) is false.

On the other hand, (*Utility*) represents the non-nihilist view, and reflects the common intuition behind what Beall calls 'standard approaches to vagueness.' The genuine distinction between bald and non-bald things is reflected in the vague predicate 'bald,' although perhaps not completely or perfectly grasped by us. The vagueness of the predicate makes us incapable of settling some borderline cases, but we are still quite able to (truly) tell a bald man from a non-bald man. Although the radical (*NoCut*) is ruled out, (*Utility*) allows for some kind of weak tolerance, e.g. if someone is clearly bald, then someone slightly hairier is still bald, though not necessarily *clearly* bald, in which case the induction step is blocked by higher order vagueness.

Beall's marriage proposal is made possible through a slight reformulation of (*Utility*), focusing on the role of the intension-part of the vague predicates. "[There] is a genuine sense in which vague predicates (their intensions) are

## 1. THE SORITES PARADOX

---

useful: they provide sharp lookalikes—sharp, homonymous relatives—of the vague” (Beall 2010, p. 188).

*(Utility')* Vague intensions yield multiple extensions that, in turn, yield sharp predicates which can be used to make true descriptions of the world.

According to Beall, this is not an alteration, but merely a slightly more elaborate expression of the original (*Utility*). It is the view that most standard approaches actually posit; the vague predicates themselves do not truly describe the world (they are vague, remember), but they stand for one or many ways of dividing the world, which we in turn use to make true descriptions of it. “Such homonyms, I suggest, are what we typically—but erroneously—call ‘vague predicates’” (Beall 2010, p. 192). The reformulated (*Utility'*) does not require that vague predicates in themselves are fit to describe the world, leaving the nihilist to do whatever she wants to them.

It remains to be seen whether this is going to be a happy marriage. I am not convinced that the nihilist is going to accept (*Utility'*), as Beall seems to think. In fact, I believe there are good reasons why she should not. If we believe that the sharp predicates yielded by a vague predicate, according to (*Utility'*), are complete and sharp (according to the nihilist’s standards) then every vague predicate must either yield every possible sharp predicate, or have imposed some kind of restriction on which predicates it yields. In the first case the multiplicity of sharp predicates is every bit as useless as the vague predicate was (and thus incompatible with (*Utility*)); drawing every possible line is no more helpful than drawing no line. In the second case the extension of the vague predicate is itself sorites-susceptible, and unless there is a special way of resolving this second-order vagueness we are caught in an infinite regress of higher-order vague predicates.

## *Precisifications of Vague Predicates*

*Achille C. Varzi (2007) points out that it is not clear how we are supposed to understand the term ‘precisification’. In this chapter I discuss how to cash out the term.*

### 2.1 VAGUE AND PRECISE PREDICATES

Predicates may be conceived as having two extensions: a positive extension, and a negative extension. The positive extension is the set of all things the predicate applies to, and the negative extension is the set of all things the predicate does not apply to. These extensions can be defined this way:

$$\varepsilon(F) := \{x \in U \mid Fx\} \quad (2.1a)$$

$$\varepsilon(\neg F) := \{x \in U \mid \neg Fx\} \quad (2.1b)$$

where  $U$  is the domain of discourse. We assume that the extensions are classical sets in the sense that they either contain an element or they do not.<sup>1</sup> Borderline cases are all the things that fails to belong in either extension of a given predicate. If Amy is too old to be clearly reckoned as a child, but still too childish to be definitely not a child, she does not belong in any of the extensions of the predicate ‘child’. In general, predicates are vague if their joint extensions do not cover the entire domain of discourse—that is, if they admit borderline cases.

---

<sup>1</sup> Some have argued that the notion of *fuzzy sets* provides a better account for the extensions of vague predicates. Membership in a fuzzy set is a matter of degree. Thus, Amy might belong to the extension of ‘child’ with strength 0.5, and to the extension of ‘not child’ with strength 0.5. (This allows vague predicates to cover the entire domain.) Although the idea has some appeal, I fail to see its contribution to the task of understanding precisifications. It might, however, be a useful framework for approximate truth under a nihilistic approach to vagueness, cf. section 1.4. For an argument for applying fuzzy sets to vague predicates see e.g. George Lakoff (1973).

## 2. PRECISIFICATIONS OF VAGUE PREDICATES

---

A defender of bivalence may object that our definition of the negative extension is formulated in such a way that it obscures the obviousness of the fact that no predicates are semantically vague. The union of  $\varepsilon(F)$  and  $\varepsilon(\neg F)$  covers the entire domain, and in a most trivial way, because the appropriate reading of ‘ $\neg F$ ’ is something like ‘... is not an element of  $\varepsilon(F)$ ’. This immediately closes the gap between the positive and negative extensions. Our reluctance to abandon the classical notion of sets secures the consequence: at the semantic level, there are no borderline cases.

Indeed, one reply is to surrender our classical notion of sets and admit that vague languages are, at the rock bottom, fuzzy. This may include predicates, relations (and hence also sets), names, quantifiers, etc., and perhaps even concepts. Fundamentally, all atomic sentences in a vague language has a degree of truth, not a strict bivalent truth-value. (Some may have truth-value 1 or 0, but this is not significantly different from having truth-value 0.472388.) Talk of classical sets is confined to a layer that simulates classicality by imposing some artificial threshold on the truth-value. In this layer we define sets of extensions  $L\text{-}\varepsilon(F)$  as e.g.  $\{x \in U \mid Fx \geq 0.7\}$  (where 0.7 is the truth-value threshold for  $Fx$ ), and we define  $L\text{-}\varepsilon(\neg F)$  as e.g.  $\{x \in U \mid Fx \leq 0.3\}$ . Although this strategy might help us escape the defender of bivalence (by rejecting her logic entirely), it commits us to an outrageous view on language (not to speak of the pile of problems that follows from rejecting classical logic). To suppose that every statement has a determinate degree of truth is no less problematic than supposing that every statement is determinately true or false. (If anything, I find the degrees of truth to be less helpful.) Unless forced otherwise, we shall not pursue this option any further.

Another reply is to grant the bivalence-defender’s reading of ‘ $\neg F$ ’. Amy is not clearly a child, which in supervaluationist terms means that she is not a child; we might express that statement this way: ‘ $\neg Child(Amy)$ ’. Hence, she belongs in the class of those who are not children, as stated by our definition of  $\varepsilon(\neg F)$ . From the fact that Amy is not a child, however, it does not follow that she is clearly not a child. Let us use the phrase ‘non-child’ for something that is clearly not a child. Amy is a borderline child, because she does not belong among the (clearly) non-children. Sometimes, but not always, she is counted in a

precisification of ‘child’; other times she is counted in a precisification of ‘non-child’. But no precisification will ever put her in both, since the positive and negative extensions are mutually exclusive. The negative extension of  $F$  is the set of things that can never, under any precisification, be  $F$ —just as the positive extension of  $F$  is the set of things that must be  $F$ . So the bivalence-defender is right, but she misses half of the picture. Our confusion is caused by reckless use of the negation operator (‘ $\neg$ ’). Without bivalence this operator fails to reflect the distinction between not being a child and being a non-child. I will try not to use it again whenever vague predicates are involved. Instead, we shall use the modified predicate ‘ $non-F$ ’ to express a condition for  $F$ ’s negative extension, which is related to  $F$  under this constraint:

$$\exists x(non-Fx \wedge Fx) \rightarrow \forall x(non-Fx \wedge Fx) \quad (2.2)$$

which says that if something is both  $non-F$  and  $F$ , then everything is both  $non-F$  and  $F$ . This implies that the predicate is meaningless, or that the antecedent is never true. If we assume that  $F$  is meaningful, which we have reason to believe is the case for many predicates, it follows that nothing is both  $F$  and  $non-F$ . (The awkward formulation is due to our recent decision not to use negation in the presence of vague predicates.) This calls for us to redefine our negative extension:

$$\varepsilon(non-F) := \{x \in U \mid non-Fx\} \quad (2.3)$$

Now we can account for the notion of precision. We say that a predicate  $F$  is more precise than  $G$  if  $F$  has more elements in its positive and negative extensions than  $G$  has in its extensions. Formally, the condition is that  $F$  is more precise than  $G$  if and only if

$$|\varepsilon(F) \cup \varepsilon(non-F)| > |\varepsilon(G) \cup \varepsilon(non-G)| \quad (2.4)$$

where  $|\varepsilon(F) \cup \varepsilon(non-F)|$  is the cardinality of the union of predicate  $F$ ’s positive and negative extension, and likewise for predicate  $G$ . This means that a predicate with 5 positive applications and 5 negative is equally precise as a predicate with 10 positive and no negative application. The only constraint is that the extensions must be drawn from the same domain. Consider the predicate ‘... has

## 2. PRECISIFICATIONS OF VAGUE PREDICATES

---

an even number of solar systems' relative to the domain of all elliptical galaxies, as opposed to the predicate '... has an even number of electrons' relative to the domain of all atoms. Because there are considerably more atoms than there are elliptical galaxies, the latter predicate is completely blown out of proportions in terms of cardinality compared to the former. But relative to a shared domain (e.g. the domain of all elliptical galaxies *and* atoms), the extensions even out, and we discover that both predicates are equally precise.

We can now also define a sharp predicate:  $F$  is sharp if and only if

$$\varepsilon(F) \cup \varepsilon(\text{non-}F) = U \quad (2.5)$$

where  $U$  is the domain of discourse. If the union of a predicate's extensions is equal to the domain, the predicate is sharp. (Specifying a domain implicitly sets the upper bound on all extensions: they cannot contain more than what is contained in the universe of discourse.)

From this, we can start sketching out a conception of precisifications. What does it mean that  $F$  is a precisification of  $G$ ? First of all,  $F$  must be more precise than  $G$ . But this is not sufficient. The predicate 'person' is by definition more precise than the predicate 'tall person' (because the former also applies to the borderline tall people), but it hardly qualifies as a precisification. 'Person' is a generalization of 'tall person', not a precisification.

Since precisifications are more precise, they must include more things. But they must also respect the extensions that are already in place. Precisifications are more precise in a particular way: a precisification  $F$  of  $G$  *extends*  $G$ .<sup>2</sup> This means that a precisification of  $G$  must build its extensions as supersets on top of  $G$ 's extensions. We use  $\supseteq$  (read: 'extends') to express this relation:

$$F \supseteq G := \varepsilon(F) \supseteq \varepsilon(G) \wedge \varepsilon(\text{non-}F) \supseteq \varepsilon(\text{non-}G) \quad (2.6)$$

' $F$  extends  $G$ ' means that the positive and negative extensions of  $F$  covers at least the positive and negative extensions of  $G$ . It is important to notice that this is a constraint on both extensions. Although 'person' is a superset of 'tall person' (it includes all tall people, and some more), 'non-person' is not a superset of 'non-tall-person' because some people are not tall, but they are still

<sup>2</sup>This term (and notation) is from Fine (1975).

people. Hence, ‘person’ does not extend ‘tall person’, and it does not qualify as a precisification—even though it is more precise (less vague).

So far, we have established the following condition on precisifications:

**Stability principle:** If  $F$  is a precisification of  $G$ , then the positive and negative extensions of  $F$  are supersets of or equal to the positive and negative extensions of  $G$ .

From this principle we can prove that precisifications cannot be less precise than the predicates they extend. But this is only a bare minimum for any viable notion of precisifications. At least for natural languages, there are further constraints on what we count as *admissible* precisifications.

## 2.2 ADMISSIBLE PRECISIFICATIONS

Apart from the stability principle, our current notion of precisifications is quite unrestricted. Nothing prevents me from making precisifications of any vague predicate, as long as they extend the vague predicate. If my use of ‘bald’ is indifferent to whether it applies to Bob, I can immediately derive two precisifications: one that counts Bob in its positive extension (‘bald<sub>1</sub>’), and another that counts him in its negative extension (‘bald<sub>2</sub>’)—as long as I remember to keep Ian McKellen in the positive, and Patrick Stewart in the negative extension of both precisifications.

But what about Al? He is slightly balder than Bob, but not definitely bald. Is he bald<sub>1</sub>? According to Fine (1975), there is a web of truths that must be retained in every precisification of a language: the penumbral truths. Some of these truths impose internal constraints on predicates; if Bob is bald, then a man with less hair than Bob is bald too. Hence, Al is bald<sub>1</sub>, since Bob is bald<sub>1</sub> and Al has less hair than Bob. Other truths restricts precisifications of predicates relative to each other, as in the Amy example. A precisification that counts Amy as a child is incompatible with a precisification that counts her as an adult, because not being an adult is part of the concept CHILD. These are external constraints. Among the penumbral truths are also the theorems of classical logic, which has given supervaluationism a strong upper hand to other non-classical semantics.

## 2. PRECISIFICATIONS OF VAGUE PREDICATES

---

Due to penumbral connections, ‘bald<sub>1</sub>’ also counts in everyone who is as bald as or balder than Bob, but it is indifferent to those who are less bald than Bob (unless they are clearly not bald). Likewise ‘bald<sub>2</sub>’ counts out everyone as bald as or less bald than Bob, but is indifferent to those who are balder. Thus ‘bald<sub>1</sub>’ and ‘bald<sub>2</sub>’ may still be vague, and we can derive further and more fine-grained precisifications from them, increasing the level of precision of our predicates. The branches can in principle be extended until every borderline case of baldness is accounted for. Although we should expect that a vast number of such precisifications may be generated (one for every thinkable way of making a term more precise), the final precisifications are sharp predicates. This means that vague predicates can be replaced by large disjunctions of sharp predicates. I call this view simple precisification:

**Simple precisification:** A vague predicate can, without altering truth, be substituted by a disjunction of all its admissible precisifications at any given level of precision.

This view entails that sentences such as ‘Bob is bald’ is equivalent to ‘Bob is bald<sub>1</sub> or bald<sub>2</sub>’, given that ‘bald<sub>1</sub>’ and ‘bald<sub>2</sub>’ cover all the precisifications of ‘bald’ at some level of precision.<sup>3</sup> If there is a ‘bald<sub>3</sub>’ and a ‘bald<sub>4</sub>’, these must be included in the disjunction as well. We may suppose that ‘bald<sub>1</sub>’ and ‘bald<sub>2</sub>’ are still vague, and that ‘bald<sub>11</sub>’, ‘bald<sub>12</sub>’, ‘bald<sub>21</sub>’, and ‘bald<sub>22</sub>’ are further and more fine-grained precisifications of ‘bald<sub>1</sub>’ and ‘bald<sub>2</sub>’, and so on. The branches of precisifications are illustrated in Figure 2.1.

Simple precisification fits perfectly with the supervaluationist conception of vagueness: a predicate is vague if there is more than one way of making it more precise, i.e. if it has more than one branch of precisifications. Here it is useful to draw a distinction between vagueness and infinity. The latter consists of predicates with one, but only one, branch of precisifications. For instance we cannot

---

<sup>3</sup>We know that ‘Bob is bald’ is a borderline case: its truth-value is *indeterminate*, viz. *true* or *false*. Furthermore, we know that ‘Bob is bald<sub>1</sub>’ is *true*, and that ‘Bob is bald<sub>2</sub>’ is *false*, and we assume that there are no more precisifications of ‘bald’ available at this level of precision, which means that the two precisifications exhaust all the ways for Bob to be bald (if Bob is bald<sub>3</sub>, he must also be either bald<sub>1</sub> or bald<sub>2</sub>, because ‘bald<sub>3</sub>’ falls under one of the former precisifications). Thus, ‘Bob is bald’ has the same truth-value, and (arguably) the same truth-conditions as the disjunction of the two precisifications: ‘Bob is bald<sub>1</sub> or bald<sub>2</sub>’.



$$\text{Bob is bald.} \left\{ \begin{array}{l} \text{Bob is bald}_1. \\ \text{Bob is bald}_2. \end{array} \right. \left\{ \begin{array}{l} \text{Bob is bald}_{1_1}. \\ \text{Bob is bald}_{1_2}. \\ \text{Bob is bald}_{2_1}. \\ \text{Bob is bald}_{2_2}. \end{array} \right.$$

Figure 2.1: Levels of precision.

give a complete (explicit) account of the extension of predicates like ‘numbers between 0 and 1’—we can always increase precision by adding another digit after the decimal point—but for every level of precision there is only one option available, and the truth-values are completely determinate at any level of precision. The predicate ‘numbers between 0 and 1’ is considered sharp because there is no more than one way of making it more precise.

Jerry A. Fodor and Ernest Lepore (1996) argue that every precisification of vague predicates is inadmissible. If our use of the predicate ‘bald’ really is indifferent to whether it applies to Bob, then, they argue, this indifference must be a conceptual truth about ‘bald’. Hence it must be *necessary* that there is no matter of fact about whether ‘bald’ applies to someone as bald as Bob, because that is part of the meaning of ‘bald’. Thus, ‘bald<sub>1</sub>’ and ‘bald<sub>2</sub>’ cannot be admissible precisifications of ‘bald’, because they violate an internal constraint on precisifications of ‘bald’. If Fodor and Lepore are right, this applies to all vague predicates; they cannot be precisified. A consequence of this is that supervenience fails. If a vague predicate cannot be completely precisified, there are no models to supervenience over. As a result, the vague predicate cannot be evaluated for truth or falsity at all.

The argument against admissible precisifications is as follows. We suppose that there is no matter of fact about whether someone one ninth of whose head is covered with hair is bald. (We call him Bob.) This means that any speaker with complete knowledge of the predicate ‘bald’, and the physical state of Bob’s head, would not be able to determine whether ‘bald’ applies to Bob. The physical state of Bob’s head is determinate, and the speaker is sufficiently competent, so any indeterminacy must depend on the *meaning* of the predicate ‘bald’. If the above is true in this world, it must be true in every world. There seems to be no possible world in which Bob is exactly as bald as he is in the actual world (and ‘bald’ has

## 2. PRECISIFICATIONS OF VAGUE PREDICATES

---

the same meaning), but in which we are somehow able to determine whether he is bald. Hence, the proposition that there is no matter of fact whether Bob is bald is not only true, but necessarily true; it is conceptually true of BALDNESS. Now, consider an arbitrary precisification  $P$  on which Bob is bald (or not bald). An admissible precisification of a model is committed to respect the penumbral truths on that model. But  $P$  flouts the conceptual truth that it is no matter of fact whether Bob is bald. Hence,  $P$  is not an admissible precisification of ‘bald’.

The crucial point here is how to interpret the statement:

- (1) There is no matter of fact about whether ‘bald’ applies to someone as bald as Bob.

According to Michael Morreau (1999), the above statement is satisfied by the fact that there are at least two complete precisifications of ‘bald’: one in which Bob is bald, and another in which he is not. This is the only way to make sense of indeterminacies, so Fodor and Lepore seem to be arguing against a necessary condition for one of their own premises. But that cannot be right.

Morreau argues that Fodor and Lepore use the phrases ‘there is no matter of fact about whether’ and ‘it is indeterminate whether’ interchangeably, suggesting that they express the same condition on sentences. Morreau applies the supervaluationist framework to account for the super-truth conditions of (1): ‘it is indeterminate whether  $P$ ’ is true if  $P$  is true on some complete precisifications and false on others. Thus, the truth of (1) presupposes that there are admissible precisifications. But Morreau’s interpretation of ‘it is indeterminate whether’ is uncharitable. If we consider the following excerpt from the original text, we see that the authors use ‘it is indeterminate whether’ in another sense than what Morreau suggests.

If it is noncontingent that there is no fact of the matter about whether someone one ninth of whose head is covered with hair is bald, then every model in which [Bob] has one ninth of his head covered with hair is ipso facto a model in which it is indeterminate whether [Bob] is bald. (Fodor and Lepore 1996, p. 523)

Fodor and Lepore claim that the indeterminacy of ‘Bob is bald’ follows from the

fact that there is no matter of fact about whether Bob is bald. The absence of a matter of fact is the primary factor in their analysis of indeterminacy. They do not presuppose a supervaluationist notion of indeterminacy; which would have made their argument against the cornerstone of supervaluationism self-contradictory. Instead, they apply a pre-theoretic understanding of indeterminacy as a third kind of semantic value (alongside *true* and *false*): the value of a statement that there is no matter of fact about whether is true or false. Surely, this is compatible with a supervaluationist understanding of indeterminacy, but does not presuppose such an understanding. Their claim is that indeterminate statements, regardless of how you conceive truth and indeterminacy, are conceptually resistant to precisification.

It might be the case that Fodor and Lepore attach a too significant role to precisifications. If we restrict our talk of colours to primaries, then the colour of a carrot might be accurately designated as the sum of yellow and red, although—strictly speaking—the carrot is neither. It has no physical properties that qualifies for falling under the concept RED nor for falling under the concept YELLOW. The property of being orange, however, is such that it has some red-like and some yellow-like qualities (as opposed to having some blue-like qualities, which it has not). In a similar way, the physical state of Bob’s head might be accurately described as bald and as non-bald, although—strictly speaking—it is neither. Such an innocuous understanding of precisifications is able to accommodate the requirement that there is no matter of fact about whether Bob is bald. Nevertheless, it is too weak for any of the precisifications to stand on their own. Just as we use the phrase ‘sum of yellow and red’ to express the same colour as ‘orange’, we can use innocuous precisifications to express the same indeterminacy as the original statement.

A more fruitful pursuit is how to interpret the modal operators in Fodor and Lepore’s argument. Suppose it is necessary that ‘Bob is bald’ is indeterminate. Here, ‘Bob’ must be interpreted rigidly; as referring to the same Bob in every modal context. (It is obvious that different Bobs differ in baldness.) Let each world represent a precisification of ‘bald’, and each accessible world (i.e. accessible from the actual—*this*—world) represent an admissible precisification of ‘bald’. This means that if ‘bald<sub>1</sub>’ is an admissible precisification of ‘bald’, then

## 2. PRECISIFICATIONS OF VAGUE PREDICATES

---

there is an accessible world in which ‘bald’ has the extensions of ‘bald<sub>1</sub>’. Fodor and Lepore’s interpretation of ‘necessary’ is that in every accessible world, the sentence ‘Bob is bald’ is indeterminate. This means that there are no accessible worlds in which Bob belongs to either the positive or negative extension of ‘bald’.<sup>4</sup> But if we interpret indeterminacy as Morreau suggests, then the indeterminacy of ‘Bob is bald’ on world  $w$  entails that there is at least one accessible world (from  $w$ ) in which ‘Bob is bald’ is true, and another in which it is false. If these worlds are also accessible from the actual world, then this is clearly inconsistent. We have claimed that the sentence ‘Bob is bald’ is indeterminate in every accessible world, but in a sense which entails that it is not. On the other hand, if these worlds are not accessible from the actual world, then, since Bob is an arbitrary borderline case of vagueness, there are no admissible precisifications. Every world in which ‘bald’ is actually precisified is inaccessible.

But there is something misguided, almost circular, about the above picture. Proponents of precisifications are forced to postulate bizarre worlds in which the sentence ‘Bob is bald’ is simply true (or false). It seems just wrong to suppose that ‘bald’ has the same meaning in such worlds as it has in ours. I suspect, however, that what Morreau has in mind with his truth condition for indeterminacy are collections of precisifications, not individual ones. He emphasizes that his truth condition is not locally valid, i.e. valid on each precisification, but *global*. “An interpretation, comprising a space of precisifications, is the sort of thing that can satisfy it. An isolated precisification is not.” (Morreau 1999, p. 154). This means that we cannot justify a modal operator that quantifies over isolated precisifications. Instead, we must consider the worlds to represent various interpretations of ‘bald’—each contains multiple precisifications. If this is true, then the claim that ‘Bob is bald’ is indeterminate in every accessible world means that every complete interpretation of ‘bald’ comprise at least one admissible precisification with Bob in its positive extension, and another with Bob in its negative extension. It is necessary that the sentence ‘Bob is bald’ is indeterminate on

---

<sup>4</sup> This framework must not be confused with the more general modal approach to vagueness. It is, however, worth noting that the accessibility-relations are strictly asymmetrical in this modal space. While ‘bald<sub>1</sub>’ is a precisification of ‘bald’, it is not the other way around. Another thing to keep in mind is that if Fodor and Lepore are right, then the modal space comprises only the actual world, or at least no other world is accessible from the actual world.

### 2.3. The relation between a vague predicate and its precisifications

---

every world, but not on every precisification.

#### 2.3 THE RELATION BETWEEN A VAGUE PREDICATE AND ITS PRECISIFICATIONS

There are two central questions concerning the nature of vagueness that arise when we consider the relationship between a vague predicate and its precisifications. The first question is whether precisifications are to be regarded as predicates in their own right, or whether they are various sharp interpretations of the vague predicate. The former perspective suggests that vagueness involves some kind of semantic overdetermination. If a term expresses a multitude of very similar, but not quite unanimous, meanings, it will be hard to tell exactly which particular meaning to focus on in any given use of the expression. This is the perspective underlying Fine's statement: "Vagueness is ambiguity on a grand and systematic scale" (1975, p. 282). On this view, precisifications are something like disambiguations. The latter perspective, on the other hand, suggests some kind of semantic under-determination: if a term does not express a sufficiently determinate meaning, there might be many possible ways of 'filling in the blanks' and gain a complete and precise understanding of its meaning. This suggests approaching vagueness in modal terms.

The second question is whether the vague predicate or its precisifications is more fundamental; is the vague predicate a product of its precisifications, or do we derive the different precisifications from the vague predicate? This question plays an important role in debates about vagueness. Suppose we knew that the vague predicate was literally defined from a cluster of precise predicates. This would in itself provide a strong argument for supervenience as it perfectly models the semantic backdrop of vague expressions. Super-truth is really the intersection of truth in a cluster of precise predicates, and thus the identification of truth and super-truth is far less objectionable.

In order to get a clearer view on the relation between vague predicates and their precisifications, let us use Peter Geach's infamous paradox of the 1,001 cats as an example. This is an instance of the more general problem of the many, which is closely related to vagueness. Tibbles is a cat (on the mat)—one single

## 2. PRECISIFICATIONS OF VAGUE PREDICATES

---

cat. Like most cats, Tibbles has at least 1,000 hairs. We use the terms  $h_1, h_2, h_3, \dots, h_{1,000}$  to refer to these hairs. Here is how Geach presents the paradox:

Now let  $c$  be the largest continuous mass of feline tissue on the mat. Then for any of our 1,000 cat-hairs, say  $h_n$ , there is a proper part  $c_n$  of  $c$  which contains precisely all of  $c$  except the hair  $h_n$ ; and every such part  $c_n$  differs in a describable way both from any other such part, say  $c_m$ , and from  $c$  as a whole. Moreover, fuzzy as the concept cat may be, it is clear that not only is  $c$  a cat, but also any part  $c_n$  is a cat:  $c_n$  would clearly be a cat were the hair  $h_n$  plucked out, and we cannot reasonably suppose that plucking out a hair generates a cat, so  $c_n$  must already have been a cat. So, contrary to our story, there was not just one cat called “Tibbles” sitting on the mat; there were at least 1,001 sitting there! (Geach 1980, §110)

We may regard the different collections of feline tissue ( $c, c_1, c_2, c_3, \dots, c_{1,000}$ ) as admissible precisifications of ‘Tibbles’. While it is not clear exactly which one of the precisifications ‘Tibbles’ refers to (hence the paradox), each precisification represents a single, definite object. According to Geach, all these precisifications are true and actual cats. Except for the fact that we think of Tibbles as one single cat, it seems reasonable to suppose that all the precisifications really are cats—they are, after all, admissible precisifications of a cat. We can ask, for instance: Is  $c_{72}$  a cat? It is hard to justify a negative answer to this question. The lump of feline tissue called ‘ $c_{72}$ ’ exhibits every intrinsic property we would expect from a cat, and few that suggests otherwise.<sup>5</sup> The same question can be asked for any  $c_n$ , until we have settled that all the precisifications are cats.<sup>6</sup> But are they cats

---

<sup>5</sup>This cat has some potentially awkward relational properties, such as the property of being a proper part of another cat, having other cats as parts, and the associated property of not being able to survive the death of another (particular) cat. These properties, however, are only awkward if we assume that the concept CAT sharply individuates cats. Apparently it does not.

<sup>6</sup>This argument assumes mereological universalism, the view that every finite collection of objects forms a composite object. I will not argue for this view here, except for briefly mentioning its alternatives: mereological nihilism, the view that there are no composite objects, or some sort of brutalism, that some (but not all) collections form composite objects. Both views solve the problem at a cost: Nihilism by denying that there is such a thing as a cat, and brutalism by more or less arbitrarily singling out one of the  $c$ -s as the only cat. (The latter may be compared to the epistemic view on vagueness, see chapter 1. Our grasp of ‘Tibbles’ must involve some

### 2.3. The relation between a vague predicate and its precisifications

---

*in their own right?* Hardly. Geach's solution to the problem is that although the *c*-s are different lumps of feline tissue, they are still the same cat. The concept CAT is vague, or 'fuzzy' as he humorously calls it, such that it is tolerant to small variations in composition. On this view the precisifications of 'Tibbles' are not 1,001 different, yet almost completely overlapping cats, but rather 1,001 sharp renderings of an initially vague cat.

We may have good and independent reasons for resisting Geach's relative notion of identity. The standard view is that identity is an absolute relation that cannot hold between different objects; not even when they differ as little as one single cat-hair (or even less, if you think that is a substantial difference). Still, Geach's solution can be helpful if we regard it not as expressing relative identity, but rather as expressing something about our concepts. The different *c*-s are not strictly speaking *the same* anything, but they still fall under the same individual concept TIBBLES THE CAT. Thus the *c*-s are the same cat in the sense that they fall under the same individual cat-concept.

Let us distinguish the following two descriptions of the present situation:

*Overlap.* Our use of 'Tibbles' denotes a multitude of almost completely overlapping objects, the *c*-s. All of them are equally entitled to being the singular referent of 'Tibbles'.

*Indecision.* Our use of 'Tibbles' provides a singular reference, but does not actually pick out any particular sharp object. The pool of possible referents includes the *c*-s.

In terms of our initial characterization of the views on precisifications, *Overlap* corresponds to the over-determination view on vagueness, and *Indecision* to the under-determination view.

In regard to the second aspect of the nature of vagueness, of whether the vague predicate or its precisifications is more fundamental, we must ask whether the fact that 'Tibbles' has precisifications makes it vague. Or perhaps it is the other way around: that we can precisify 'Tibbles' because the term is vague. The answer to this question will help us decide which of the above descriptions that is most accurate. One might say that the outcome is pretty much the same:

---

margin of error because we lack perfect knowledge of the concept CAT.)

## 2. PRECISIFICATIONS OF VAGUE PREDICATES

---

‘Tibbles’ is both vague and has precisifications. Still, this generalizes to our understanding of precisifications and the nature of vagueness, and would provide good motivation for either accepting or rejecting a procrustean approach to vagueness.

The problem of the many extends into higher-order vagueness. The condition for falling under the concept TIBBLES THE CAT is just as impossible to provide as the condition for ‘being the same cat as  $x$ ’. Although we might solve the problem with Geach’s 1,001 particular cats, there are millions of other ways to ‘carve’ up Tibbles; many of which convey objects that are not as clearly cats as the  $c$ -s. The difficulty might not be to provide various precisifications, but to draw the line between admissible and inadmissible precisifications.

### 2.4 CENTIPEDE GAMES AND UNSHARPENABLE VAGUENESS

John Collins and Varzi (2002) provides an example of vague predicates that arguably cannot be precisified. Consider the following version of the game called ‘Centipede’: There are two players,  $X$  and  $Y$ , and a pile of \$200 in silver dollars (each coin is \$1) on the table. The players take turns drawing either one or two coins from the pile until there is no money left on the table. They keep whatever they draw. If a player draws two coins, the game ends instantly and the remaining pile of money is removed.<sup>7</sup> The diagram in Figure 2.2 illustrates the course of this game. If the players cooperate throughout the game, they end up with \$100 each.




Figure 2.2: Diagram of a Centipede game.

A horizontal move (*across*) leads to the next round of the game, slightly increasing the total pay-offs for both players, and a vertical move (*down*) ends

<sup>7</sup>This version of the game is inspired by John Broome and Wlodek Rabinowicz 1999.



## 2.4. Centipede games and unsharpenable vagueness

---

the game and the players walk away with their respective earnings. The potential pay-offs are displayed as a pair under each vertical line; player  $X$ 's pay-off to the left and player  $Y$ 's to the right. In any mode, the game is considered a *leave-it game* if the logical consequences of moving across are more favourable for the current player than the logical consequences of moving down. If it is the other way around the game is a *take-it game*. We assume that both players aim only to maximize their own pay-off (they have no interest in their opponent's gain) and that they are rational throughout the game. The players are rational to the extent that they believe they are able to tell whether the game is a take-it game or a leave-it game, form beliefs and act accordingly; if the rational player believes that the game is a take-it game, she also believes that she should move down. Both players also believe that their opponent is rational (this is crucial for the forthcoming argument).

At the end of the game, when there are only \$2 left in the pile, the game is clearly a take-it game. Player  $X$  should of course take both coins, because she is rational (and  $\$101 > \$100$ ). But this determines the rational choice in the previous round as well. When there are \$3 on the table, player  $Y$  should move down and gain a total of \$100. Player  $Y$  knows that if she plays across, player  $X$  will draw two coins in the next round (as we have already established that she will). This will end the game and leave player  $Y$  with only \$99, as opposed to the \$100 she will get by drawing two coins at this point. So in this round the game is also a take-it game, because player  $Y$  is rational and she believes that player  $X$  is rational as well. This, in turn, determines the further previous round the same way; player  $X$  moves down, anticipating player  $Y$ 's next move, and so on. By backward induction, we can determine every rational move in the game this way. The conclusion is that in every mode of the game—even the initial—the game is a take-it game. Therefore, if both players are rational (in the above sense), and believe that their opponent is also rational, the first player will end the game in its first round, taking the two dollars.

There is something absurd, almost paradoxical, about this picture. In order to maximize her winnings, player  $X$  settles for \$2 in a game that could have earned her fifty times more. At the same time she knows that if the game continues past the next round, she will already have gained more than \$2—and will continue to

## 2. PRECISIFICATIONS OF VAGUE PREDICATES

---

gain as long as the game moves along. It is in the interest of both players to keep the game running almost as close to the end as possible. A cooperative strategy appears to be the most rewarding. Each player knows that, although she certainly runs the risk of being cut one dollar short from her theoretically maximum gain in the final rounds of the game, she will nevertheless gain significantly from keeping the game running, as opposed to ending it as quickly as possible. She also believes that the same goes for her opponent, and that the other player is aware of this fact as well. (Remember, these players aim only at maximizing their own gain—it is not a goal to ‘beat’ the opponent by making more money than her.) Thus she expects her opponent to reason similarly, and the game is clearly a leave-it game.

Let us sum up the situation. We have the following two predicates, defined for a game  $g$  in round  $n$  where  $i$  is the last round:

**Take-it game:**  $g_n$  is a take-it game iff  $n + 1 \geq i$  (i.e. this or the next round is the last, so moving down is most favourable).

**Leave-it game:**  $g_n$  is a leave-it game iff  $n + 1 < i$  (i.e. this or the next round is not the last, so moving across is most favourable—you are confident that the game will not end before your next turn).

The problem is that for any  $n$ , if  $g_n$  is a take-it game, then  $n$  is effectively the last round (i.e.  $n = i$ )—provided that both players are rational, etc.—because the current player is rationally obliged to end the game whenever it is a take-it game. From this it follows that if  $g_n$  is a take-it game, then so is the previous  $g_{n-1}$ . On the other hand, for any  $n$ , if  $g_n$  is a leave-it game, then the next round cannot be the last, nor can the one after that, and so on. Assuming that  $i$  is a natural number greater than 0, under these conditions a game can either have one single round, or infinitely many. In the first case, the single game is a take-it game, and in the last case, every game is a leave-it game.

This reveals that there is no paradox lurking behind this Centipede game. A contradiction comes only from assuming both that the game is finite, and that there is at least one round in which the game is a leave-it game (as Collins and Varzi seem to do). But the single-game solution is unsatisfactory in many ways. On a larger scale, the logical consequences of moving across seem to out-favour

## 2.4. Centipede games and unsharpenable vagueness

---

the logical consequences of moving down early. There is an ambiguity of scope in ‘rational obligation’; if we focus on single cases we are rationally obliged to move down, but if we consider a series of cases we are rationally obliged to move across.

Thus we have not yet provided a satisfactory definition for ‘take-it game’ and ‘leave-it game’. It is a conceptual constraint on both predicates that they apply (with certainty) only if the current player is rational and believes she will gain most from acting upon it (i.e. moving down on a take-it game, and across on a leave-it game). But we do not know how to decide this.

Collins and Varzi suggest that the predicates ‘take-it game’ and ‘leave-it game’ are vague. They claim that in the final round, when there are only two coins left on the table, the game is as clearly a take-it game as it is a leave-it game in the first round (where all 200 coins are present).<sup>8</sup> The backward induction argument for Centipede is really a sorites that exploits the blurred line between take-it games and leave-it games.

The vagueness of these rationality predicates is peculiar. Unlike the kind of vagueness we find in observational predicates, the rationality kind of vagueness appears to resist any attempt of making it sharper; it does not seem to have any admissible precisifications at all. Contrast with a precisification of, say, the predicate ‘tall’ (we call it ‘tall#’) such that everyone taller than 170 cm is tall# and everything else is not. ‘Tall#’ is an admissible precisification of ‘tall’ because it observes the penumbral connections of ‘tall’: nothing non-tall is counted into the positive extension of ‘tall#’, nothing is both tall# and non-tall#, etc. It also meets a further condition, something Collins and Varzi calls ‘public accessibility’: it is possible, in principle, for speakers to shift their linguistic practice such that the meaning of the precisification ‘tall#’ replaces the original meaning of ‘tall’ in their language. Whatever vagueness is left after this shift, can be accounted for by epistemic disability and a margin of error principle—usually, we are unable to tell with absolute precision whether an approximately 170 cm tall person is taller than 170 cm or not.

---

<sup>8</sup>Actually, their claim is somewhat stronger. They claim that in the final round the game clearly *is* a take-it game, and that in the first round the game clearly *is* a leave-it game. Since they neither argue for this (strongest) view, nor need it for present purposes, I allow myself to settle on a slightly weaker formulation. I suspect that there are no clear cases in this game.

## 2. PRECISIFICATIONS OF VAGUE PREDICATES

---

The first condition is pretty straightforward. In order to pass as an admissible precisification the new predicate must inherit the conceptual constraints (penumbral connections) of the related vague predicate. Otherwise we are not speaking of an actual precisification, but of a similar, yet distinct lexical entity.

The second condition is somewhat harder to justify. If a precisification is a more or less arbitrary decision to chop off semantic vagueness by picking one of many possible sharper meanings of a word, why should we require it to be publicly accessible which precisification we have picked? This requirement seems impossible to meet beyond a certain level of precision, for instance when two precisifications of 'bald' differs by one single hair. On the other hand, we may argue that this is not impossible at all, at least not in principle. Given sufficient time, cognitive abilities and observational powers (which it does not seem completely outrageous to suppose that human beings have obtained in some possible world or distant future) we may overcome this obstacle. But that hardly justifies public accessibility as a condition for admissible precisifications.

Collins and Varzi argue that if every speaker was allowed to precisify vague predicates in her own idiolect, in such a way that it was impossible for other speakers to know the details of her private precisifications, then we would not obtain any precisifications at all. We would instead have millions of homonyms; similar, but not equal, predicates for 'bald', making the expression ambiguous on a grand scale. Alternatively, we may insist that all the private precisifications are candidate meanings of the same predicate. It follows that the vagueness of the original predicate is preserved in the totality of precisifications used by the speakers. Although the vagueness of this 'new' predicate is caused by some awkward statistical variation in use, it is nevertheless just as vague as the original predicate was (given a sufficient number of competent speakers, time, etc.). Unless we assume that there is one true precisification for every vague predicate, such that all other precisifications are more or less right relative to it, we are forced to accept the condition of public accessibility. Without accessibility, precisifications can never emerge in any language governed by a community of speakers.

The predicates 'take-it game' and 'leave-it game' are unsharpenable because they cannot have precisifications that fulfil both conditions. If player *X* precisi-

## 2.4. Centipede games and unsharpenable vagueness

---

fies ‘take-it game’ to include  $g_n$ , and this precisification is accessible to player  $Y$ , by backwards induction player  $Y$  will precisify *her* notion of ‘take-it game’ to include  $g_{n-1}$ , and so on.

We have so far not discussed whether these predicates are actually vague. It might be the case that they are not vague at all, and that their apparently strange behaviour can be accounted for in terms of some kind of shortcoming in rationality, rather than a deficiency in truth or meaning—after all, the centipede game itself does not depend on our descriptions of it. On the other hand, the game is essentially about judging what is the most favourable move. This directly involves predicating to the game whether it is a take-it game or a leave-it game. Without such a dimension, the game is just a meaningless exercise. If there is vagueness involved in this game, then it is clearly an obstacle for procrustean approaches to vagueness. Unlike observational predicates, ‘take-it game’ and ‘leave-it game’ do not appear to have any ‘blurred’ edges or borderline areas. Any arbitrary division is as (un)reasonable as the other. But by making a cut, both predicates collapse into the same default position: every game is a take-it game. Precisifications seem to provide no help to us here.



## Chapter Three

# *Frege on Sharpness*

*This chapter discusses Frege's sharpness requirement as a response to vagueness. We will argue that the scope of this requirement is restricted to what we may call a 'logical point of view', and that Frege, in light of Russell's paradox, regards a set-theoretical understanding of his notion 'extension of a concept' as defective.*

### 3.1 THE SHARPNESS REQUIREMENT

Frege considers logic to be concerned with sharp concepts only. He requires that every concept must be completely defined, i.e. clearly stipulated for any object whether it falls under the concept or not.<sup>1</sup> Constructions that do not meet this requirement are not considered to be concepts; hence, by definition, there are no vague (or otherwise incomplete) concepts. This principle is often referred to as Frege's sharpness requirement.

**Sharpness Requirement:** Every concept must have sharp boundaries such that it is unambiguously determinable of every object whether it falls under the concept or not.

$$\forall F \forall x (Fx \vee \neg Fx) \tag{3.1}$$

where the second-order variable  $F$  is a concept.

This view is commonly attributed to Frege, and formulations of it seems to appear frequently in his writings. The principle derives from the law of excluded

---

<sup>1</sup>See Frege 1903, p. 69.

### 3. FREGE ON SHARPNESS

---

middle, and has therefore a firm position within classical logic. It is our interest in classical truth that forces this restriction upon us.

If it is a question of the truth of something—and truth is the goal of logic—we also have to inquire after *Bedeutungen*; [...] we have to throw aside concept words that do not have a *Bedeutung*. These are not such as, say, contain a contradiction—for there is nothing at all wrong with a concept’s being empty—but such as have vague boundaries. It must be determinate for every object whether it falls under a concept or not; a concept word which does not meet this requirement on its *Bedeutung* is *bedeutungslos*. (Frege 1969, p. 133)

The ‘concept word’ is a predicate in Frege’s terms. The predicate’s denotation, its *Bedeutung*, is a concept.<sup>2</sup> Frege argues that the predicates that fail to meet the sharpness requirement have no *Bedeutung*; such predicates are *bedeutungslos*—they do not successfully denote any concept, in the strict Fregean sense of ‘concept’. Such predicates are attempting to name something that in fact does not exist. Frege frequently speaks of such terms as ‘mythical’ or ‘fictional’.<sup>3</sup> They behave as if they were meaningful predicates, but they are not, just like names of fictional characters—like ‘Luke Skywalker’—behave as if they are meaningful names. But ‘Luke Skywalker’ does not pick out any person, because there is no such person to pick out.<sup>4</sup> Likewise, the predicate ‘bald’ does not pick out any

---

<sup>2</sup>Frege’s use of ‘concept’ is slightly different from the common use of this term. According to Frege, a concept is a special kind of function that maps every object to a truth-value: the True, or the False. If an object is mapped to the True by a given concept, then that is an object falling under the concept. We may speak of the set of all objects falling under a concept as its extension. For present purposes, most differences between Frege’s particular use of ‘concept’ and the common use of the same word are unimportant.

<sup>3</sup>Cf. Frege 1903, p. 76; Frege 1969, pp. 133–4; Frege 1979, pp. 269–270.

<sup>4</sup>Frege solves the lack of *Bedeutung* for proper names by introducing a function  $\backslash\xi$  that basically maps the argument onto itself; if  $\Delta$  is a singular object (viz.  $\Delta$  is identical to the extension of the concept IDENTICAL WITH  $\Delta$ ), then the value of  $\backslash\Delta$  is the object  $\Delta$ , otherwise the value of  $\backslash\Delta$  is the collection of none or many objects denoted by  $\Delta$ . Thus, although the name ‘Luke Skywalker’ might not have a *Bedeutung*,

$\backslash(\text{extension of the concept THE PERSON NAMED ‘LUKE SKYWALKER’})$

does. Cf. Michael Beaney 1997, pp. 384–385 and *Grundgesetze, I*, §11. This solution is not available for predicates, since they can only have concepts as their *Bedeutung*, and the function



concept, because there is no Fregean concept that would satisfy the vague sense of ‘bald’. Fictional terms have a sense, but no reference.

In a letter to the mathematician Guisepe Peano in 1896, Frege turns to this observation to reject the sorites paradoxes as falsidical: “The fallacy known as ‘Acervus’ [Latin for ‘heap’] rests on this, that words like ‘heap’ are treated as if they designated a sharply delimited concept whereas this is not the case” (Gabriel et al. 1980, p. 114). The sorites argument is a falsidical paradox, and the fallacy is to depend on vague predicates in reasoning about truth. In the argument we use the predicate ‘heap’ in an attempt to say something assertive, and thus act as if it was completely defined. Therefore, we expect every instance of ‘ $Fx$ ’ to come out as true or false. This is obviously not the case, as the sorites arguments demonstrate. Due to their incomplete definition, predicates like ‘heap’ have only a *Sinn* and not a *Bedeutung*.<sup>5</sup> We understand the sentence “This pile of four beans is a heap”, but has no way of determining whether it is true or false, because ‘heap’ is not defined for that case.

The sharpness requirement seems to have the consequence of rendering vague predicates ultimately meaningless.<sup>6</sup> As long as there are any potential borderline cases, the vague predicate does not name a sharp concept. Therefore, it does not name anything.

(*NoRef*) Vague predicates fail to meet the Sharpness Requirement and thus have no *Bedeutung*.

Gary Kemp (1996) argues that (*NoRef*) commits Frege to a form of semantic nihilism for ordinary languages. The reason is that Frege is already committed to the view that a sentence  $Fa$  is true only if the individual referred to by ‘ $a$ ’ falls under the concept designated by ‘ $F$ ’. But there are few, if any, predicates

---

<sup>5</sup>Frege 1969, pp. 133–134.

<sup>6</sup>Frege’s notion of ‘vague’ includes gappy predicates, for instance

$$\begin{aligned} n \text{ is nice* if } n \leq 3 \\ n \text{ is not nice* if } n \geq 4 \end{aligned}$$

Here the predicate ‘nice\*’ is not defined for values between 3 and 4, but it is not vague in the relevant sense. This predicate involves a truth-value gap, but it is not sorites-susceptible. Frege’s broader notion of ‘vague’ is not important here, and we will restrict our discussion to the relevant sense of the word.

### 3. FREGE ON SHARPNESS

---

in ordinary language that actually meet the strict criteria of the Sharpness Requirement. Consequently, most predicates in ordinary discourse should fail to name anything. If  $F$  is vague, it follows that the predicate ‘ $F$ ’ does not refer to any concept for the object  $a$  to fall under (regardless of what the value of ‘ $a$ ’ is). Hence most sentences uttered in ordinary language do not yield a truth-value, and as a result “... precious little of what we actually think and say is true or false” (Kemp 1996, p. 172).

According to Kemp, Frege never addresses this issue, as if he is either unaware of it, or accepts it but does not consider it important.<sup>7</sup> Others claim that (*NoRef*) is falsely attributed to Frege. Stephen Puryear (2013) argues that Frege imposes a sharpness constraint on predicates only from what he calls ‘a logical point of view’: the pursuit of a strict and classical truth-functional language. Although authors like Kemp are unsympathetic to the idea that Frege might be considering other points of view besides the logical, the textual evidence for Frege’s sharpness requirement seem to involve a more or less explicit emphasis on the logical point of view.<sup>8</sup> In “Comments on Sinn and Bedeutung” Frege seems to provide straightforward evidence for (*NoRef*).<sup>9</sup> But closer inspection reveals that the entire claim is wrapped in a conditional statement: “If it is a question of the truth of something—and truth is the goal of logic ...”. This suggests, according to Puryear, that Frege thinks that there is (at least) another point of view from which vague predicates are regarded as having a *Bedeutung*—although that *Bedeutung* is probably not fit to be labelled a ‘concept’ according to Frege’s standards.

A similar strand is found in Joan Weiner’s (2007) interpretation of Frege. Although she, unlike Puryear, accepts that Frege holds (*NoRef*), she denies that Kemp is right in submitting this as a problem for Frege. Weiner points out that Frege, within the scope of his logicist project, carefully distinguishes regarding a sentence as being true from claiming that a sentence actually *is* true. “Although Frege writes as if the terms of everyday language have *Bedeutung* and the sentences of everyday language have truth-values, he never actually *says* that

---

<sup>7</sup>Braun and Sider (2007) suggest the latter, without arguing for their interpretation of Frege.

<sup>8</sup>See Puryear (2013) for textual references and discussion.

<sup>9</sup>The relevant paragraph is quoted earlier in this section, see p. 40 above.

they do.” (Weiner 2007, p. 706). She suggests that Frege is aware of the consequences identified by Kemp, but he does not consider it problematic to deprive ordinary language most of its apparent truth-values. On the contrary, we should never have expected to find classical validity and truth in vague languages, such as ordinary language. Frege’s aim is to prove the basic truths of arithmetic, and his texts must be read with this limited scope in mind. The *Begriffsschrift* is not intended to be an ideal language, but rather a tool designed for a very specific purpose: “We have need of a system of symbols from which every ambiguity has been banned and whose rigorous logical form cannot be escaped by the content.” (Frege 1964, p. 158). From this point of view it is neither unfortunate nor ruinous (as Kemp characterizes it) that most, if not all, expressions in ordinary discourse—even ordinary arithmetic discourse—fails to observe the strict, truth-functional criteria of the *Begriffsschrift*. In fact, it seems quite irrelevant. Well aware of the shortcomings of validity in ordinary languages, Frege wisely refrains from asserting any sentence expressed in such a language without reservations.<sup>10</sup>

Weiner suggests that Frege might be operating with two different notions of truth: one strict and classical notion of truth for the *Begriffsschrift*; and some other notion that does not require a *Bedeutung* for ordinary languages. She compares the latter notion to the supervaluationist notion of truth. Some expressions in ordinary language might indeed express truths, in the sense that they are true on every precisification of the expression.

After all, the significance of Frege’s regarding it as true that each number has unique successor is that on every acceptable definition (precisification) of the term ‘number’, it will be provable, hence true,

---

<sup>10</sup>According to Frege, whenever we use an expression we presuppose that it has a *Bedeutung*.

Now we can of course be mistaken in the presupposition, and such mistakes have indeed occurred. But the question whether the presupposition is perhaps always mistaken need not be answered here; in order to justify speaking of the *Bedeutung* of a sign, it is enough, at first, to point out our intention in speaking or thinking. (We must then add the reservation: provided such a *Bedeutung* exists.) (Frege 1892, pp. 31–32)

### 3. FREGE ON SHARPNESS

---

that each number has a unique successor. (Weiner 2007, pp. 709–710)

Frege wants to provide a scientific language with precise definitions that will prevent any intrusion of vagueness and ambiguity from ordinary language. A definition of ‘number’ in this language will single out one complete precisification of the term. If Frege is right in regarding statements involving the term ‘number’ in ordinary language as true, this precisification is to be found among the actual precisifications of our ordinary, incomplete, notion of ‘number’. But there are important differences between Frege’s project and the supervaluationist project. Supervaluationists consider an expression to have its meaning fixed in true sentences, such that a precise definition of ‘number’ cannot cross the boundaries of the already existing notion. Whatever come out true under the ordinary notion, must also come out true under the precise notion. On the other hand, Frege insists that, until a complete and sharp definition of the concept has been provided, we cannot know whether there is any concept to speak of at all—we are forced to work under the assumption that such a concept exists. If our pre-systematic notion has failed to capture a particular concept, a complete definition can never be provided, and everything based upon the assumed semi-concept must be reconsidered. “[A concept] is something we try to grasp, and hope in the end to have grasped correctly, if we have not mistakenly looked for something where there was nothing to find.” (Frege 1971, pp. 204–205).

#### 3.2 AXIOM V

So far, we have not discussed our grasp of the notion ‘extension of the concept  $F$ ’. In chapter 2, we defined the extension of a predicate as the set of things that the predicate applies to. According to Frege, a predicate names a concept, so a predicate applied to the name of an object is an expression of the case that the object falls under the concept. So we can understand the extension of a concept as the set of things falling under that concept. For instance, let us assume that there is a (sharp) concept RED which is successfully named in our language by the predicate ‘red’. Then there are some things that satisfy the condition ‘ $x$  is

red’ (or perhaps none—it does not matter, there is nothing wrong with an empty extension). These things comprise an object in their own right, which we call the extension of the concept RED—or the set of red things.<sup>11</sup>

Frege is committed to the above picture by his infamous Axiom V, which serves to help us recognize extensions of concepts as logical objects. If two concepts are such that the same objects always falls under them, then they have the same extension. This is not a definition, but rather a recognition that having all and only the same objects falling under two concepts *is the same* as these concepts having the same extension.<sup>12</sup>

**Axiom V (for concepts):** The extension of the concept  $F$  is the same as the extension of the concept  $G$  if and only if the concepts  $F$  and  $G$  are materially equivalent.

$$\varepsilon(F) = \varepsilon(G) \leftrightarrow \forall x(Fx \leftrightarrow Gx) \quad (3.2)$$

This axiom entails the existence of an extension to every concept in Frege’s system, which is what Frege needs in order to define numbers in terms of extensions.<sup>13</sup> The extensions of concepts are conceived as sets, and thus logical objects.

Logicians have long since spoken of the extension of a concept, and mathematicians have used the terms set, class, manifold; what lies behind this is a similar transformation; for we may well suppose that what mathematicians call a set (etc.) is nothing other than an extension of a concept, even if they have not always been clearly aware of this. (Frege 1903, p. 278)

<sup>11</sup> Remember that concepts, according to Frege, are governed by the Sharpness requirement. Thus, it follows that any  $x$  either falls under the concept denoted by ‘is red’ or it does not. Therefore the conclusion that the red things form a set seems pretty straightforward: if you have some things, then you have the set of those things. (In fact, here is that set:  $\{x \mid x \text{ is red}\}$ .) For a defence of the principle of plural collapse (the collapse of pluralities into sets), see Øystein Linnebo (2010).

<sup>12</sup> Readers familiar with Frege will notice that what we are considering here is a special case of Axiom V. The general case covers any kind of function, not just concepts. (It can be attained by replacing the words ‘extension’ and ‘concept’ with the words ‘value-range’ and ‘function’, respectively.) But this distinction is not important for our current discussion.

<sup>13</sup> See Richard G. Heck (2011) for a discussion of why Frege wanted to define numbers in term of extensions.

RUSSELL'S PARADOX

There is a well known inconsistency in Frege's system, known as Russell's paradox. Since the extensions of concepts are logical objects, they are themselves among the things that can fall under their own concept—and thus belong to themselves. This is not immediately a problem. But we may consider the extension of the concept of a set that does not belong to itself: does it belong to itself? If it does, then it is not a set that does not belong to itself (and hence it does not belong to itself). If it does not belong to itself, then it *is* a set that does not belong to itself, and hence it does belong to itself. But that is a contradiction!

This is one way to generate Russell's paradox. When Axiom V is introduced we can derive the principle of Naïve Comprehension:

**Naïve Comprehension:** For any concept, there is a set of all and only the things that falls under that concept.

$$\forall F \exists y \forall x (x \in y \leftrightarrow Fx) \quad (3.3)$$

This principle is naïve in the sense that it imposes no restrictions on the kind of concepts it quantifies over. Unfortunately, its lack of restrictions is precisely what gives rise to a version of Russell's paradox. Consider the concept of a set that does not belong to itself. According to Naïve Comprehension, its extension is the set of all and only the sets that are not members of themselves.  $[\lambda z . z \notin z]$  is the concept of a non-self-membered set, which we instantiate for  $F$  in (3.3):

$$\exists y \forall x (x \in y \leftrightarrow [\lambda z . z \notin z]x) \quad (3.4)$$

$\lambda$ -conversion yields

$$\exists y \forall x (x \in y \leftrightarrow x \notin x) \quad (3.5)$$

We suppose that  $r$  is the *Russell set* (i.e. the set of all non-self-membered sets), and instantiates the existential quantified variable  $y$  with  $r$ .

$$\forall x (x \in r \leftrightarrow x \notin x) \quad (3.6)$$

The final step is to ask ourselves: 'Is the set of all non-self-membered sets a member of itself?' We can do this by instantiating the universally quantified

variable  $x$  with  $r$ . This gives

$$r \in r \leftrightarrow r \notin r \quad (3.7)$$

which is an explicit contradiction. The Russell set is a member of itself if and only if it is not a member of itself. If we suppose that it is, then it is not; and if we suppose that it is not, then it is.

We shall not discuss Russell's paradox in too much detail, but there is one important point that is relevant to our understanding of Frege's response to vagueness. Stephen Yablo (2004) points out that Naïve Comprehension is really the combination of two distinct principles (p. 151):

**Naïve Plurality Comprehension:** For any concept  $F$ , there are some things that are all and only the things that fall under  $F$ .

**Naïve Set Comprehension:** Any plurality of things forms the set of those things.

In order to block Russell's paradox, it is sufficient to reject one of these two principles. While the standard approach has been to reject Naïve Set Comprehension by imposing some kind of limitation on what kinds of sets that can exist, Yablo argues for rejecting Naïve Plurality Comprehension.

Nevertheless, Naïve Set Comprehension seems to be the source of our commitment to the existence of the Russell set. At least at first glance. If we consider all the non-self-membered sets, Naïve Set Comprehension tells us that there must be a set of those things—the Russell set. But the first step in this reasoning presupposes Naïve Plurality Comprehension: we assume that we can consider all the non-self-membered sets, i.e. all and only the sets that fall under the concept of not belonging to themselves. Perhaps we cannot.

The property  $P$  that (I say) fails to define a plurality can be a perfectly determinate one; for any object  $x$ , it is a determinate matter whether  $x$  has  $P$  or lacks it. How then can it fail to be a determinate matter what are all the things that have  $P$ ? I see only one answer to this. Determinacy of the  $P$ -s follows from

- (i) determinacy of  $P$  in connection with particular candidates,

(ii) determinacy of the pool of candidates.

If the difficulty is not with (i), it must be with (ii). (Yablo 2004, pp. 151–152)

Frege is committed to Sharpness, so we cannot expect to find any object for which it is indeterminate whether it falls under a given concept—which corresponds to Yablo’s talk of properties: the property of falling under a given concept. Thus, if we follow this diagnosis, the potential members of the Russell set are drawn from an indeterminate pool of candidates. Michael Dummett (1991) suggests that Frege’s biggest mistake is that he does not have “the glimmering of a suspicion of the existence of indefinitely extensible concepts” (p. 317). Ordinals and sets are examples of indefinitely extensible concepts; it is provable that every ordinal ( $n$ ) has a successor ( $n + 1$ ), and (although this is slightly more controversial) that there is no universal set of ‘everything’: it is in the nature of sets to combine and build up indefinitely, so every set is member of a bigger set.<sup>14</sup> If this is correct, we are wrong to suppose that there are some sets that are all and only the non-self-membered sets; not because the concept of a non-self-membered set is flawed or otherwise incoherent, but because there is no way to assess for this property in *every* set.

This is perhaps as close as we get to a *vague object* in Frege’s ontology: a set with a sharp and well-defined condition for membership, but with an indeterminate pool of candidates. The size of this set varies under different realizations of its pool of candidates. In the unqualified state that the Russell set appears in the above argument, it has no size, in the sense that it has no *determinate* size (and Sharpness forbids us to recognize predicates that do not apply determinately). Apparently, the Sharpness requirement in itself is not a sufficient shield against vagueness. Even if it can be determined for every object whether it falls under a given concept, there is no guarantee that the concept has determinate extensions; which is presupposed in Axiom V. What is missing is a determinate domain of quantification.

---

<sup>14</sup>For any set  $x$ , we can construct the *power-set* of  $x$ ; it has as its members every possible set formed by members of  $x$ . The power-set of  $x$  is itself a set, but is different from  $x$ , and it too has its own power-set, etc. This process goes on forever, i.e. indefinitely.



Patricia A. Blanchette (2012) argues that it is unreasonable to suppose that Frege's quantifiers are not restricted to any determinate domain, even though Frege himself is neither clear nor explicit on this point. "From the Fregean point of view, in which sentences are fully interpreted and theories are sets of thoughts, there is always a determinate subject matter under discussion that already sets the boundaries of quantification" (p. 74). When Frege speaks of arithmetic of integers, the quantifiers are appropriately restricted so that ' $\forall x$ ' means 'every integer'. If this applies to the quantifiers in Axiom V, then Naïve Comprehension and Russell's paradox does not follow. But nor do the extensions of concepts as determinate logical objects.

In the appendix to the second volume of *Grundgesetze*, where Frege acknowledges the recent discovery of Russell's paradox, he admits "that the interpretation we have so far put on the words 'extension of a concept' needs to be corrected" (Frege 1903, p. 256). Although he never completes the task, he seems to recommend that we abandon the idea that the extension of a concept, in the sense of a set, yields absolutely all objects that fall under the concept.<sup>15</sup> This might be a recognition of the distinction between what Linnebo (2010) describes as the intensional nature of concepts and the extensional nature of sets. While sets are identified in terms of their elements, concepts are identified in terms of their condition for application. The concept of a stone is what it is because it only applies to stones. The set of all stones, however, is what it is because it comprises the very elements that it does. In a world where there are twice as many stones as in our world, the concept STONE remain the same, but its extension, the set of all stones, will be a different one. Therefore, the left-to-right direction of Axiom V must be false; we cannot derive the fact that concept  $F$  and  $G$  are materially equivalent from the fact that their extensions (when construed as sets) coincide.

---

<sup>15</sup>Frege later speaks of Russell's paradox as dealing "the death blow to set theory itself" (Frege 1979, p. 269). By 'set theory' he means the kind of naïve set theory that follows from Axiom V, where sets are comprehended in terms of concepts.

### 3.3 FREGE ON ORDINARY LANGUAGE

Frege distinguishes between ordinary language and his own *Begriffsschrift*. He describes them as having different features and purposes, without rejecting the former as defective. In the preface to *Begriffsschrift* he clearly recognizes the utility of vagueness and imprecision in ordinary language.

I believe that I can best make the relationship of my *Begriffsschrift* to ordinary language clearest if I compare it to that of the microscope to the eye. The latter, due to the range of its applicability, due to the flexibility with which it is able to adapt to the most diverse circumstances, has a great superiority to the microscope. Considered as an optical instrument, it admittedly reveals many imperfections, which usually remain unnoticed only because of its intimate connection with our mental life. But as soon as scientific purposes place great demands on sharpness of resolution, the eye turns out to be inadequate. The microscope, on the other hand, is perfectly suited for just such purposes, but precisely because of this is useless for all others. (Frege 1879, p. V)

The claim that a logically perfect language is too precise and therefore useless for any purpose beside the scientific is quite remarkable. Frege clearly appreciates the flexibility of vague language in ordinary discourse, which suggests that he does not think vague predicates are deficient in meaning. It is the procrustean approach of scientific language that cuts off the flexible parts of vague predicates through its inability of recognizing non-rigid entities.

We may interpret this as a claim that ordinary discourse does not require strict truth and falsity, as we discussed in section 3.1. Just as a casual glance at someone is insufficient to determine precisely how many hairs that person has on his head, vague predicates like ‘bald’ are insufficient to provide determinate truth-values. But ordinary discourse requires a certain level of coarseness; it has no use of too fine-grained distinctions. This suggests a supervaluationist notion of truth; if the sentence ‘Bob is bald’ expresses something true, although neither ‘Bob’ nor ‘bald’ actually refer to something, it is because any admissible precisi-

fication of ‘Bob’ (that can be recognized by science) falls under any admissible precisification of ‘bald’ (that can be recognized by science). But this is a long shot.

In fact there is little reason to suppose that the Sharpness requirement is intended as a prerequisite for truth. Blanchette (2012) points out that Sharpness serves to secure the *preservation* of truth in deductions. Logically perfect languages must be maximally truth-preserving, such that any kind proof conducted within the system is valid. Such languages cannot recognize predicates that may fail to return a truth-value for some argument. Nevertheless, while this makes it clear that ‘heap’ cannot be recognized in a scientific language, since it would potentially generate a sorites paradox, this gives us no reason to deny that sentences involving the word ‘heap’ can express something truth-evaluable in language in general.

The proof theoretic justification for the Sharpness requirement is supported in one of Frege’s letters to Peano:

It can be objected [against the Sharpness requirement] that [vague] words are used thousands of times in the language of life. Yes; but our vernacular languages are also not made for conducting proofs. And it is precisely the defects that spring from this that have been my main reason for setting up a conceptual notation. (Gabriel et al. 1980, p. 115)

Nowhere does he claim that the ‘defects’ of ordinary language consist in its lack of truth or meaning. The shortcoming is rather its lack of strict validity. But even if not all inferences are valid in ordinary language, we are nonetheless able to express true (or false) thoughts. In fact, Frege is quite explicit about this:

The task of our vernacular languages is essentially fulfilled if people engaged in communication with one another connect the same thought, or approximately the same thought, with the same proposition. For it is not at all necessary that the individual words should have a sense and meaning of their own, provided only that the whole proposition has a sense. (p. 115)

### 3. FREGE ON SHARPNESS

---

It is clear that Frege is talking about ordinary language as expressing thoughts, and we have no indication that these should be any different from the kind of thoughts expressed in a logically perfect language. Thus there is no reason why they should be judged true or false in another sense of ‘true’ and ‘false’.<sup>16</sup>

Williamson (1994) criticizes Frege’s account of vagueness. Vague sentences are treated like idioms; their meaning is not determined by the individual words, but by the entire sentence. For example, the phrase ‘kick the bucket’ does not get its meaning from its compound words (there is usually neither kicking nor buckets involved), but from the phrase itself. A learner of English who has not yet learned the meaning of this exact phrase cannot be expected to understand it, even she knows all the grammar and individual words involved.

The claim that sentences with vague words should function in a similar way is highly implausible, considering the amount of successful communication in ordinary discourse. If we only understand a vague sentence in itself, and not by means of its compound words, we should not understand sentences we have never encountered before. But, unlike the English learner and the phrase ‘kick the bucket’, as competent speakers we can understand ‘a heap of bald men’ even if we have no previous experience with that particular phrase; which strongly suggests that we have some understanding of its compound words.

Perhaps Williamson puts too much weight on Frege’s claim that the individual words are not required to have a sense and reference of their own. The idiom-interpretation, where *no* word has a (relevant) meaning on its own, is certainly the strongest possible reading of the claim. But if we continue to read Frege’s letter, he seems to aim at a more modest and plausible claim.

Where inferences are to be drawn [...] it is essential that the same expression should occur in two propositions and should have exactly the same meaning in both cases. It must therefore have a meaning of its own, independent of the other parts of the proposition. In the case of incompletely defined concept words [i.e. vague predicates], there is no such independence: what matters in such a case is whether

---

<sup>16</sup> Of course, a sentence in ordinary language may fail to express a proposition at all, or convey something totally different from what the speaker intended—as opposed to in formal languages, where a unique thought is guaranteed to follow from any well-formed formula.

the case at hand is the one to which the definition refers, and that depends on the other parts of the proposition. (Gabriel et al. 1980, p. 115)

Again, Frege is in remarkably clear opposition to the view that vague predicates are meaningless. His point is that vague predicates have no *independent* meaning, in the sense of a unique Fregean concept; but that is not necessary for a complete proposition to make sense. As a semi-concept, 'bald' is perfectly meaningful in the sentence 'Patrick Stewart is bald', since 'bald' is defined for Patrick Stewart. But perhaps it has no meaning in 'Bob is bald', if 'bald' is not properly defined for Bob. Thus, whether 'bald' has meaning or not in a particular sentence depends on whether it is defined for the case at hand.<sup>17</sup>

While Frege might not be as dismissive of vagueness in general as he appears to be in his main works, it is important to notice that Frege never actually develop semantics for ordinary language. We do not know whether there is a deeper insight to his remarks on ordinary language. What we do know, however, is that the pressure for a procrustean approach to vague predicates comes from 'the logical point of view'; here it is not sufficient that only some sentences express a thought—it is required of every sentence. And such a requirement is too rigid for vague predicates.

---

<sup>17</sup> Another feature that seems to worry Frege with improperly defined predicates is their more general context dependence. For instance, the predicate 'tall' seems to have a completely different meaning in the sentence 'Tally is tall' compared to the sentence 'The Empire State Building is tall'. In the last sense of 'tall', Tally is definitely not tall. But this is not directly vagueness-related.



## Chapter Four

# *The Many-boundaries Approach*

*Mark Sainsbury (2013) criticizes the experimental set-up in traditional sorites arguments. Our inclination towards the principle of tolerance can be explained in terms of other factors, and is not a unique feature of vagueness. Sainsbury's own proposal is discussed.*

### 4.1 THE FORCED MARCH PHENOMENA

A conditional sorites argument is usually presented in terms of the following (or a similar) thought experiment. Suppose we have a soritical series of colour patches ranging from red to orange. There is a sufficient number of patches such that the difference in shade between one patch and its adjacent is not significant in terms of colours. I show you the first (clearly red) patch, and ask “Is this patch red?” You answer “Yes.” I switch to the next patch, and ask “Is this patch red?” You answer “Yes.” I switch to the next patch, ask the same question, and you proceed by answering to the best of your ability. And so on, and so forth. For each step you are forced to provide one single answer; it might be “yes”, “no”, “don’t know”, “perhaps”, etc. How many kinds of answers you are allowed to make is unimportant. What matter is that you answer appropriately and to the best of your ability for each step, and that you are consistent in your way of answering. If you answer “No” in one case, then you are most likely going to answer “No” in the next case, since there is (as far as you can tell) no relevant difference between the two cases—and you have already committed your answer to one of them. This kind of experimental set-up is called a ‘forced march’.<sup>1</sup> In a forced march you will always have access to your most recent

---

<sup>1</sup>This formulation of a forced march is from Priest (2003).

#### 4. THE MANY-BOUNDARIES APPROACH

---

previous answers and take them into account when you move up (or down) the sorites series. You can compare the last case to the present, and see (because this is a sorites series) that the difference between the two cases is insignificant, if they are distinguishable at all. The result is that for each step, any shift in semantic value seems unreasonable.

According to Sainsbury (2013) we can expect the following behaviour from the subject of a forced march:

*Overextension.* The subject will apply the initial predicate (“Yes, that patch is red.”) to more cases than she would have done under normal circumstances.

*Sudden Change.* The subject will eventually flip; changing her answer to “No” (or “Don’t know”, or anything else), although there is no relevant (or even distinguishable) difference between the two adjacent cases.

*Hysteresis.* The point where the subject flips is sensitive to which direction the sorites series is presented to her. (The experiment could be reversed; starting with the clearly orange patch.)

*Anxiety.* The subject will experience anxiety and frustration from not knowing how to answer correctly in the borderline area.

In order to explain this behaviour we can turn to some version of the Tolerance principle. This will yield a ‘semantic’ explanation of the forced march phenomena: scoreboard contextualism. Alternatively, we can turn to the Margin of Error, and provide an ‘epistemic’ explanation. Sainsbury argues in favour of the latter.

#### SCOREBOARD CONTEXTUALISM

Contextualists can explain what goes on in a forced march by applying a super-valuationist or open-texture approach to the meaning of the vague predicate in question. When the parties in a conversation accept a certain sentence as true, e.g. ‘This (particular item) is red’, that sentence is put on the ‘conversational scoreboard’. When the next case is presented, the contextual commitments on



the scoreboard must be taken into account. Whether the present case is red, is determined by our accepted use of ‘red’ in the current conversational context.

The scoreboard straightforwardly explains Overextension and Hysteresis. The previous answers are already up on the board, and Tolerance secures that the predicate applies in the same way to cases that do not differ significantly. Since we suppose there is no significant difference between the present case and the previous, the subject is right in extending her former judgement—whatever that might have been, depending on the direction of the sorites series. This will, presumably, result in the subject overextending her initial application of the predicate, which explains both phenomena.

Sudden Change is explained by the fact that, eventually, the conversational scoreboard will appear incoherent to the subject. When the subject is looking at a clearly orange patch—and ‘realizes’ that the patches must have been orange for quite some time—the extension of ‘red’ is stretched beyond its permissible use. In this case, the context, which also includes common beliefs about redness, will assign both *true* and *false* to the sentence ‘This patch is red’. Such a context is (clearly) incoherent.<sup>2</sup> The subject will dismiss the current context, clear the conversational scoreboard, and start over by giving a new and (presumably) different answer to the present case. Thus, a Sudden Change has occurred.

Anxiety is the most difficult phenomenon to explain. It will certainly happen when there is tension between the conversational scoreboard and Tolerance on one hand, and the subject’s impression of permissible use of the predicate on the other. (In other words, in the vicinity of Sudden Change.) As the context becomes increasingly incoherent, the level of anxiety will rise. Sainsbury, however, predicts that “we can reasonably expect uncertainty and anxiety to set in earlier [than the point where the context is significantly incoherent], and this is what contextualists lack the resources to explain” (2013, p. 231). The reason for this early anxiety is that the subject anticipates what is coming. Yet we have no way to explain this behaviour since the contextual adjustment of the predicate’s

---

<sup>2</sup>Hans Kamp (1981) describes contextual coherence as a vague concept: “A context can become gradually less and less coherent but it is impossible to say for sure precisely when it is coherent no longer” (p. 253). I am sure he is right regarding coherence as an epistemic notion. But I suspect that, for present purposes, an understanding along the semantic notion of consistency is sufficient. We do not need another layer of vagueness.

#### 4. THE MANY-BOUNDARIES APPROACH

---

extension is supposed to make sentences true—although the very same sentence might be borderline or false in other contexts.

##### EPISTEMIC EXPLANATION OF THE PHENOMENA

The epistemic explanation is compatible with the view that there is an objective and sharp meaning to every predicate, or at least that there might be—for all we know. The forced march phenomena are explained by our ignorance of such precise facts of predicates, from which it follows that we have no justification for making subtle distinctions. Sainsbury (2013) describes the epistemic reasoning this way:

I have just judged  $\#n$  to be red and I can detect no relevant difference between  $\#n$  and  $\#n + 1$ . So [...] I should not discriminate these patches, that is, I should judge  $\#n + 1$  to be red also. (p. 230)

This explains Overextension and Hysteresis. Reasoning as above will get your red pretty far across the orange side.

Still, our judgements are primarily based on what we see and know. Seeing a clearly red patch, we will judge it to be red. But sooner or later the subject to a forced march will be presented with a patch that she can see is clearly orange, and change her answer. This is a Sudden Change. Since there is no relevant difference between that patch and the previous, the subject will probably want to change her mind about that as well (and so on).

Anxiety follows naturally from the subject's experience of uncertainty. As the forced march proceeds into the borderline area, the subject is basing most of her justification on her previous judgements; the subject matter that is presented to her is insufficient for an independent judgement. The subject will probably anticipate this as soon as she understands the pattern of the experiment, whether she is familiar with slippery-slopes arguments or not. Therefore we can expect Anxiety even before the really uncertain cases are presented.

Compared to scoreboard contextualism, the epistemic account provides a more robust explanation of the forced march phenomena. While contextualism presupposes a particular semantic theory, the epistemic account is compatible

## 4.2. Forced marches reveal nothing interesting about vagueness

---

with any kind of semantics—as long as we do not *know* exactly how the predicate works. Furthermore, the epistemic account has a ready explanation of Anxiety, which scoreboard contextualism lacks the vocabulary to explain. The last point, that we will discuss further in the following section, is that it turns out that forced marches do not depend on vagueness. If a forced march can be constructed from precise terms, then this cannot be explained by scoreboard contextualism, who depends on some level of open texture.

If we accept the epistemic account as the best explanation of forced marches, then we cannot expect to learn anything about vagueness from forced marches. The epistemic account is compatible with the view that there is no vagueness, and is expressed in terms of judgements and justification.

### 4.2 FORCED MARCHES REVEAL NOTHING INTERESTING ABOUT VAGUENESS

Sainsbury (2013) provides the following version of a forced march:

Guided by a casual glance around the stadium you are in, you are asked to give your best estimate of how many people are there, at a minimum. Your estimate is correct if there are at least that number of people. You respond “3,000”. Your interlocutor says: “Given the coarse-grained character of your evidence, isn’t the answer 3,001 just as likely to be right?” It’s hard to suppose that you could rationally think that 3,001 would be more likely to be false than 3,000, for, if you retain your belief that 3,000 is a correct lower limit, the supposition that there are not at least 3,001 people entails that there are exactly 3,000 people, and you know you are in no position to make a sensible estimate of exactly how many people there are. In light of these considerations, you agree that 3,001 is as good an estimate. Under these pressures, a series of answers will emerge: 3,000, 3,001, 3,002 . . . and so on, until you reach a breaking point, and decline to slide further down the slippery slope. (p. 232)

#### 4. THE MANY-BOUNDARIES APPROACH

---

In this example we suppose that it is not vague how many people are at the stadium, hence it is not vague to which numbers the predicate ‘there are at least  $n$  people’ applies. Yet the subject is caught in what clearly resembles a forced march, and is likely to experience the forced march phenomena. Only this time without any vagueness whatsoever.

It might be objected that, although the number of people in the stadium is not a vague matter, your estimate might be. After all, what the interlocutor is calling into doubt is not the actual number of people at the stadium, but the accuracy of your estimate. Your estimate is imprecise in its nature. You are asked to give an approximate number of people based on a quick glance around the stadium. Of course that is inaccurate! It would be purely accidental if you should happen to guess the actual number of people at the stadium. You do not possess the kind of information required to tell exactly how many people there are. Therefore, any estimate concerning the number of people must be regarded as a rough approximation—clearly within the Margin of Error.

This objection can be met. Suppose that, instead of asking for your estimate, the interlocutor shows you a number of photographs of the stadium with various crowds. For every photograph there is a definite number of people present (and amazingly, they are all visible on the photograph). You are asked to point at a picture with no more than the number of people at the actual stadium. As before, the actual number of people is not a vague matter, nor is the predicate ‘no more than the actual number of people’. But this time we are confident that there is a definite number of people in your “answer” as well. Suppose you point at the photograph with 3,000 people. If your interlocutor asks you whether the picture with 3,001 people is just as likely to be a correct answer, you obviously cannot tell the two crowds apart and must admit that it is. And so on and so forth. Once again you are caught in a forced march.

This kind of forced march cannot be explained by scoreboard contextualism. The precise meaning of every crucial term is fixed, so there is no room for contextual shifts in the extensions of anything. We must apply to your lack of knowledge to explain this forced march. First of all, you have no precise knowledge about the number of people at the stadium. You do, however, have a rough approximation based on your glance around the stadium. Thus, you have

## 4.2. Forced marches reveal nothing interesting about vagueness

---

what Williamson (1994) calls ‘inexact knowledge’;<sup>3</sup> you have *some* knowledge about the number of people present, but not enough to know the exact number. The imprecise nature of this knowledge gives you no reason to prefer answering 3,000 people over 3,001 people. If you suppose that there are more than 3,000 people, you must also suppose that there are more than 3,001 people. But you still know that there are no less than, say, about 500 people and no more than, say, about 10,000 people (the accuracy will probably depend on your previous experience with estimating crowds).

The kind of reasoning involved in a forced march, and hence also in the conditional sorites, has nothing to do with vagueness in particular. The crucial factor is the memory of your previous answers, which serves as a base for justification of similar cases. If you were unable to access and compare your previous commitments to the present case, you would (arguably) not display Overextension or Hysteresis.

Sainsbury (2013) argues that “If our aim is to obtain the best estimates we can of the minimum number of people in a stadium, we would never allow the forced march procedure. This makes it puzzling why one would take forced march sorites as the canonical guide to the nature of vagueness” (p. 232). It is the forced march procedure itself, and not the subject matter, that distorts the results and produces the forced march phenomena. These phenomena do not depend on vagueness in particular, but on a specific kind of epistemic state.<sup>4</sup> Therefore, we should not expect forced marches to reveal anything whatsoever apart from our subject’s epistemic attitude towards it.

I agree with Sainsbury that vagueness is not an epistemic phenomenon. Therefore it is surprising that he seems to overlook an important difference between the examples of the red patches and the stadium. Suppose, on your brief glance around the stadium you notice the big screen that says: “There are 3,751 people here right now”. (And let us suppose that this particular stadium is known for its accurate counting of visitors.) If you believe that the number on the screen is accurate, you will be able to respond confidently, without any risk of being

---

<sup>3</sup>Cf. section 1.3.

<sup>4</sup>This will not bite on those who consider vagueness to *be* this kind of epistemic state with regard to meaning, e.g. when you have only a loose grasp of the precise meaning of ‘red’.

#### 4. THE MANY-BOUNDARIES APPROACH

---

caught in a forced march. But no similar information can be provided for the red patches example. Even if you knew everything you could know about redness, the human eye, light, the physical configurations of the various patches, and so on, you would still not be able to sharply tell the red from the orange patches—not because you would not observe or know the precise shades of the patches, but because that information would not help you determine precisely which of the patches that are red.

Is it true that forced marches reveal nothing interesting about vagueness? Priest (2003) claims that the key to understanding the vagueness puzzle is the concept of a cut-off—which stands out as the core issue in a forced march sorites. But the forced march seems only to unveil that there is such a difficulty; whether it is due to vagueness or ignorance (or other causes) is not indicated. For instance the concept of a person might turn out to rest largely on our lack of knowledge.<sup>5</sup> Suppose we had accurate, true, and reliable information about exactly under what conditions the potential (or actuality, depending on your preferred beliefs) of consciousness and other relevant mental states begins and ceases to exist in humans. In this case, the idea that a dying person might stop existing from one particular moment to the next seems far less outrageous than if the concept of a person is really vague (or worse: if personhood is metaphysically vague). I see no kind of information that can enlighten us in a similar way regarding redness, heaps, baldness, and so on. Put differently: If we were to learn a true way to sharply delimit red from non-red, and adjust our use of the predicate ‘red’ accordingly, we will have changed the lexical meaning of ‘red’. On the other hand, if we were to learn the exact number of people at the stadium, the exact size of the universe, or how to precisely discern a dying person from a corpse, we would change our knowledge and epistemic toolbox—not the meaning of the

---

<sup>5</sup> Two remarks: (1) There are many concepts of a person—some of which are truly vague—but I encourage the reader to focus on a concept that might turn out to have determinate boundaries. (Call it a ‘scientific person’, if it helps.) This does not mean that personhood must begin and cease abruptly. Just like a film clip might fade out at the end, our personhood might fade out when we die. In the film clip, however, it is not vague where the fade out-transition begins. There is a determinate point in the film where the editor starts the process of fading out the clip.

(2) I do not claim that all facts are knowable—not even in principle. Hence the difference between some kinds of forced marches might never be revealed. I do not argue for this view here.

words. In order to respect the notion of semantic vagueness we should not even accept the possibility of a true and precise meaning of predicates like ‘red’.

In conclusion, a forced march is indifferent to various sources of ignorance. As a test it reveals lack of precise knowledge—not vagueness. The vagueness-related forced marches might be explained in terms of scoreboard contextualism, but they might as well be explained in epistemic terms. There is no glorious argument here, but the fact that the epistemic explanation covers both vagueness and non-vagueness forced marches contributes to its advantage. On the other hand, a proponent of semantic vagueness should resist a purely epistemic explanation of the vagueness-related forced marches and instead seek a different approach that enables her to rule out such an explanation.

### 4.3 SCRAMBLED SORITES

In order to avoid the forced march phenomena, Sainsbury suggests that we use a different experimental set-up: The ‘scrambled sorites’. In a scrambled sorites, cases are presented to the subject in a random order, not in the soritical order. Once the subject has given her answer, she is distracted in some way to ensure that she will not remember the specifics about her previous answers. Suppose I show you a patch and ask you whether it is red. Once you have replied, I remove the patch, play loud music, flash the lights, ask you to build a house of cards, or whatever it takes to distract your attention. Now you should have no clear memory of exactly what shade of red you have ‘committed’ to. When I show you a different patch and ask whether it is red, your reply will not depend on your previous judgements—at least not as much as in the forced march.

Once you have given your answer to every case, we can rearrange them to correspond to the soritical order of the patches, such that your response to the most clearly red patch is to the left, and your response to the most clearly orange patch is to the right. If your pattern of answering is consistent, in the sense that you do not switch back and forth between different replies, your score is a ‘perfect score’. For example, if you have only answered “Yes” and “No”, your score should consist of a series of “Yes”s followed by a series of “No”s. If you answered “Don’t know” to some patches, the series of “Yes”s should be followed

#### 4. THE MANY-BOUNDARIES APPROACH

---

by all the “Don’t know”s before the series of “No”s. (And so on for even more nuanced replies.)<sup>6</sup>

Now, the assumption is that there are many perfect scores associated with a scrambled sorites. While a perfect score in itself is a remarkable result, a situation in which every perfect score is the same—that is, scores that make every cut at precisely the same patches—is unbelievable. It is plausible that, given a sufficient number of trials, the scrambled sorites will yield distinct perfect scores (let us call them  $S_1$ ,  $S_2$ ,  $S_3$ , and so on). The different scores assign redness and non-redness differently to the patches, such that a given patch might be red according to  $S_1$ , but not according to  $S_2$  and  $S_3$ .

Distinct perfect scores are mutually exclusive, since they assign opposite truth to the same sentences. Therefore, at most one of them can be true. If we suppose that it is the case that e.g.  $S_3$  is true, we will take something like an epistemicist attitude towards the vagueness of ‘red’: We assume that there is an objective meaning of the predicate ‘red’, but that we have no way of knowing exactly what that meaning is.<sup>7</sup> Otherwise, if we do not wish to accept such semantic mysticism (we do not, cf. section 1.3), we must accept that there is a sense in which no possible response to a scrambled sorites is true.

Sainsbury suggests that the variations between perfect scores reflect different *permissible*, but not mandatory, uses of the vague predicate. You are allowed to make one out of many possible cut-offs between the red and non-red patches—as long as your use is within the constraints of the applied concept. These constraints include in particular the requirement that if patch # $n$  is red then any redder patch is also red, and the constraint that some things are clearly red while

---

<sup>6</sup> In the original experiment, as described in Sainsbury (2013), the subject is required to answer strictly “Yes” or “No” to every case—there are no other options. But he seems to admit further kinds of responses, under the condition that they can be arranged in an orderly structure with the most definite “Yes” on one side and the most definite “No” on the other. While the number of options probably has a significant impact on the results, it makes no difference to our a priori considerations. A perfect score must involve a consistent pattern of answering and sharp cut-offs, regardless of how many.

<sup>7</sup> Supervaluationists who do not take higher-order vagueness seriously, are also committed to this kind of attitude. The true score is the one that says “Yes” to all the clearly red patches, “No” to all the clearly non-red patches, and “Don’t know” or “Indefinite” to the others. This score assigns redness to the patches that are red under all precisifications of ‘red’, and non-redness to the patches that are non-red under all precisifications. But due to higher-order vagueness there is no matter of fact which patches are red under all precisifications.



other things are clearly not. Any arbitrary division that respects those constraints is permissible, but there is no such division in particular that is mandatory. In this sense, no perfect score is strictly speaking true. For example, the manager of a paint shop is allowed to display tubes of red paint on one shelf, and tubes of orange paint on another, despite the fact that the reddest orange is virtually indistinguishable from the most orange red. If the shelves or the store's selection of tubes had been different, the manager could have chosen another way to divide red and orange—without risk of self-contradiction or incoherence.

Vague predicates are associated with many boundaries, yet no one in particular. Every boundary corresponds to a possible perfect score. Boundaries are, in this sense, precisifications that are related to the vague predicate in the sense that they are possible and admissible interpretations of it. What separates Sainsbury's account from our previous candidates is that none of these precisifications are true in any appropriate sense. For example, what counts as a safe distance is radically different when you are watching a tennis match and when you are watching an airplane take off. But these different interpretations are not true in virtue of the meaning of 'safe distance'. They are locally justified in virtue of their context, which includes elements of risk in the particular kinds of scenarios and some arbitrary choice to make it enforceable. Similarly, the division of red and orange tubes of paint is not true in any sense of the word 'red', but it is justified by the manager's wish to display these particular tubes on the two shelves in an ordered manner. If there is an odd number of tubes, the manager is free to determine whether the middle tube is going on the red or the orange shelf—both cuts are permitted under the meaning of 'red'.

How does a scrambled sorites hold up against the stadium example? If I ask you, for any random number (within a reasonable range) whether there are more people in the stadium than that number, and take steps to ensure that you do not remember your answer once you have replied, we will eventually have a similar kind of score. But since there is no vagueness in this matter, there *is* one (and only one) perfect score that is true. For any other score, your answer is wrong on at least one occasion.

Should we accept Sainsbury's generous account of vague predicates? He certainly provides valuable insight by considering vagueness from a fresh per-

#### 4. THE MANY-BOUNDARIES APPROACH

---

spective, but his notion of boundaries suffers from many of the same shortcomings as the similar notion of precisifications. If we suppose that boundaries are sharp, in the sense that they provide a condition for dividing things into two (or more) groups, then these boundaries *are* some kind of precisifications—defined extensionally by perfect scores. The concept of a perfect score is indefinitely extensible, unless you want to impose some arbitrary restriction on the samples in your sorites series (as we admittedly have done in all the examples so far). Thus, a perfect score is merely possible, and can never actually be provided. Although this does not prevent us from conceiving boundaries, the above characterization of vagueness turns out to be no more familiar than the notion of admissibility or fuzzy sets. All in all, Sainsbury is most eager to stress that many cut-off points are permissible for vague predicates. But that conclusion does not follow from the scrambled sorites, nor from the notion of a perfect score.

Sainsbury's most remarkable contribution to the vagueness debate is his usually gentle approach to the phenomenon. Vagueness is not considered a hard, technical problem that must be resolved within a strict framework of logic. Rather, it is a soft and flexible feature of thought and language, that cannot be forced to fit the existing paradigms of logic and set theory. In an earlier paper, he writes:

Anything worthy of the name *boundary* will effect set-theoretically describable divisions, even if more complex ones than the simple twofold division envisaged just now. But any such division, however complex, will misdescribe the functioning of a vague predicate.  
(Sainsbury 1996, p. 253)

In Sainsbury (2013) we are urged to interpret the above statement as the claim that vague predicates has no *unique* boundary. In at least one sense, many boundaries are not collectively worthy of the name 'boundary'. Still, any collection of boundaries are subject to supervaluation and therefore yields something boundary-like. But that is mere imprecision, not vagueness. The difference in plurality comprehension is important; the phrase 'many boundaries' does not necessarily denote something definite, as 'collection of boundaries' does.<sup>8</sup> If

---

<sup>8</sup> Cf. section 3.2.

Sainsbury stands by his claim that “any such division, however complex, will misdescribe the functioning of a vague predicate”, it must be his view that there are indefinitely many such boundaries, which prevent supervaluation of the vague predicate in itself—only partitions of the vague predicate may be supervaluated.

In conclusion, Sainsbury’s many-boundaries approach to vagueness lack specification at a most crucial term: the ‘many’ in ‘many boundaries’. If this is understood extensionally, as referring to the many boundaries that are or might be drawn, the many-boundaries approach harmonizes with the view that vague predicates make sense only through their precisifications. On the other hand, we may understand ‘many’ as referring to some boundaries, but none in particular. On this view, the vague predicate makes as much sense in itself as it does under any precisification. This is more in line with Sainsbury’s previous writings, and his statement that “it is the vague predicate just as it is that invites the many cuts, potentially registered as the supervaluationist’s sharpenings. They are not products of some process of revision, but simply reflect permissible uses” (2013, p. 235). If this is right, then Sainsbury, disguised as Procrustes, will invite you to spend the night in any bed from his vast collection of different sized iron beds; or, if you would prefer, you might as well sleep on the couch.

#### 4.4 QUANTITIES

Our characterization of vagueness so far depends on the view that some properties ‘come in degrees’. This means that we assume there are properties that tell us how tall, how bald, how red (etc.) something is. We may call these properties ‘quantities’. We are able to order things according to their quantities, for instance their height or redness. This is a prerequisite for soritical series and for the application of vague predicates to individuals (e.g. whether you are tall depends on your height).

But the quantitative aspect of properties is not captured in language. Quantity is completely absent in predicate logic, and it seems to have a poor foundation in other languages too. Consider height. In order to describe my height as a property in English, I say something like ‘I am 185 cm tall’ or perhaps ‘I have

#### 4. THE MANY-BOUNDARIES APPROACH

---

the height 185 cm'. I do not understand this as me having some generic property of height to the degree 185 cm. Rather, I have the property that is my particular height: the property of being 185 cm tall. Thus, my height and yours (unless we are the same height) are two different properties.<sup>9</sup> We have a firm understanding of height as an intrinsic property, i.e. a property that you possess independently of anything outside yourself. We also have a strong inclination towards accepting height as a quantity: a property that comes in degrees. It makes sense to order things according to height, and to determine whether something is taller than another. This, however, requires that two properties (my height and your height) are the same in one sense, and different in another. In respect to height, you and I are the same in the sense that we both have height, but distinct in the sense that we have different heights.

Quantitative properties are described by Lewis (1986) as “families of plain properties” (p. 53) and as relations to numbers. These are obviously two different things, and correspond to the two different aspects of height. I have an intrinsic property that is my height, and a height-in-centimetres relation to a real number close to 185. According to this view, we do not order things height-wise based on their intrinsic properties, but based on the number to which they have a height-in-centimetres relation to. (For convenience I will continue to speak of such relations as properties.)

In the case of height, the relation to a number is given by a unit of measurement. But this does not cover all kinds of quantitative properties. Not all properties that come in degrees can be reduced by measurement to relations to numbers. Such reduction requires a scale and the possibility of measuring properties relative to it, at least in principle. And, apparently, not all properties can be measured. Consider for instance complex phenomenal properties like pain or excitement. There is no doubt that we experience various levels of excitement in different states of mind, but I think it is impossible to measure these levels on any scale—regardless of how many dimensions it might have—without a significant loss of accuracy.<sup>10</sup> Another example, perhaps more in the spirit of this paper,

---

<sup>9</sup>That is, different *kinds* of properties. I do not mean different in the trivial sense that two distinct and particular instances of something are not identical.

<sup>10</sup>This is not an empirical problem. I am sure that psychology will eventually develop impres-

is the relevance of utterances in a conversational context. Different replies may be considered more or less relevant at any given point in a conversation. Still, it seems impossible to accurately measure this level of relevance.

Even though the degrees of some properties cannot be measured, we seem to have something like a concept of excitement and relevance, that enables us to tell whether one reply is more or less relevant than another in a given context—provided that they are discernible in relevance.<sup>11</sup> This enables us to order things according to relevance. Suppose we have a conversational context and five potential utterances ( $u_1, u_2 \dots, u_5$ ), that differ in relevance. We suppose that each  $u$  is discernible (in the relevant sense) from some of the other  $u$ -s, but not necessarily all. Now we can exploit the schema ‘In the current conversational context,  $u_n$  is more relevant than  $u_m$ ’ to determine which sentences that are true of the  $u$ -s.<sup>12</sup> If there is an order of the  $u$ -s that respects all those sentences, then we have ordered the  $u$ -s according to relevance. It might be the case that we are unable to determine whether  $u_4$  or  $u_5$  is most relevant. Both may be more relevant than the other  $u$ -s, but indistinguishable from each other. Extending our model to include more utterances might help us discover a more fine-grained order. There might be a  $u_6$  that is more relevant than  $u_4$ , but indistinguishable from  $u_5$ . This will settle the order of  $u_4$  and  $u_5$ . Finally, it is likely that many  $u$ -s will be having the same relevance, and their internal order will therefore be arbitrary. But this is a fact that we can only learn inductively; by never discovering any difference in relevance between them.

The outcome is not a precisification, nor anything like it. We have not divided the  $u$ -s into any kind of groups. What we have done is to describe a complex

---

sively accurate and sophisticated means for measuring and scoring various complex psychological states. I anticipate, however, that some relevant information will not survive the translation into a unit of measurement.

<sup>11</sup>If I ask: “Is Malmö the capital of Sweden?”, replies like “Yes”, “No”, “I don’t know”, etc. are equally relevant. In comparison “I’ve never been to Sweden” is less relevant, and “42” is completely irrelevant—at least apparently (you never know).

<sup>12</sup>It might be objected that the notion of relevance is ambiguous (even in this narrow context). There are different ways in which the relevance of two utterances may vary, so the schema does not secure that we are speaking of relevance in the exact same sense across all sentences. I think this can be met by adapting some rigid convention. For instance, after the first sentence has been evaluated, all subsequent sentences refer to the kind of relevance applied in the evaluation of that first sentence.

#### 4. THE MANY-BOUNDARIES APPROACH

---

relation between all the *u*-s, that serves as a conceptual guide for our notions of ‘more’ and ‘less’ of some property.

##### 4.5 OBSERVATION AND RISK

We have used something I called ‘concept’ to individuate so-called properties that come in degrees. But this is a very different kind of concept than what Frege called ‘concept’: a function that maps arguments to truth-values.<sup>13</sup> It depends on our ability to distinguish between different cases, or perhaps more accurately it *is* our ability to distinguish between cases with respect to some kind of feature. Wright (1976) argues that this capacity reflects the semantics of observational predicates, which makes it impossible that they are governed by a Fregean concept or supervaluationist interpretation. “[If] we are able to make such a distinction, there can be no objection to introducing a predicate to express it. And then, it seems, the semantics of *this* predicate will have to be observational” (p. 243). The semantics of such predicates are then more in tune with some kind of psychological notion of concepts: learned collections of traits that we use in order to categorize various things. And indeed, even our actual grasp of a Fregean concept is a concept in the psychological sense; over time we have learned to place things in different categories, and we believe that some of these correspond to predicates. But psychological concepts are subjective, and language is not. The sentence ‘Patrick Stewart is bald’ is true—regardless of who utters it. It is not a claim that Patrick Stewart falls under my subjective concept of baldness, but a claim that he really *is* bald, or at least that I believe he is. Nevertheless, it is our individual abilities to distinguish bald and non-bald people that determines whether it is true. In the case of Bob, our ability is of little help, and, since there are no inaccessible proper meaning of ‘bald’, there is no objective answer to whether ‘Bob is bald’ is true; the community of competent English speakers are unable to categorize Bob as bald in a way that is consistent with our usual distinction between bald people and non-bald people.

As we learned from Frege, the problem with vague predicates is not that they are unable to figure in true sentences. The problem is that once they are allowed

---

<sup>13</sup> Cf. chapter 3.

into a language, that language is no longer strictly logical. In a logically perfect language we can trust every inference to preserve truth, and therefore be safe from paradoxes. This does not mean that every inference in ordinary language is false or paradoxical. According to Saul A. Kripke (1975) we must allow there to be an element of risk in our semantic theories.

[An] adequate theory must allow our statements involving the notion of truth to be *risky*: they risk being paradoxical if the empirical facts are extremely (and unexpectedly) unfavorable. There can be no syntactic or semantic “sieve” that will winnow out the “bad” cases while preserving the “good” ones. (p. 692)

Although Kripke has the Liar paradox in mind,<sup>14</sup> the lesson we must learn is that we take a risk when we admit vague predicates in logical reasoning. Frege’s unaccomplished aim of a logically perfect language is one without any risk. Kripke, on the other hand, wants to include a truth-predicate, and he is aware that this makes his language vulnerable to semantic paradoxes under some extremely unfavourable conditions. Including vague or observational predicates entails an even greater risk that some inferences are paradoxical, but it does not invalidate the entire language. The ability to express obvious truths might warrant taking the risk, and the speaker must be careful not to abuse it on unexpectedly unfavourable cases.

So whether precisifications should be part of a logic of vagueness depends on the level of risk we are willing to take. The cost of not taking the risk is some kind of linguistic impairment. We have seen that the analysis of vague predicates in terms of sets of precisifications fails to reflect the meaning those predicates have in ordinary language; we do not actually use precisifications to determine whether something is red, someone is bald, or whether I am in the outback. We use casual observation and cognitive abilities to distinguish between cases; when we are unable to make a distinction, even if we understand that there is or might be a difference, we do not distinguish between the cases.

---

<sup>14</sup>The Liar paradox is briefly as follows: is the sentence ‘This sentence is false’ true? If it is, then it is not; and if it is not, then it is.

#### 4. THE MANY-BOUNDARIES APPROACH

---

On the other hand, it seems possible to imitate most features of vagueness. So the question is really whether we want our theory to include vagueness or not. If we want a logic that includes vague predicates, we must either accept that there are paradoxical ‘black holes’ in that logic (and try to steer clear of them), or let Procrustes do his thing.



## *Bibliography*

- Barnes, Jonathan (1982). “Medicine, Experience and Logic”. In: *Science and Speculation*. Cambridge University Press.
- Beall, Jc (2010). “Vague Intensions: A Modest Marriage Proposal”. In: *Cuts and Clouds*. Ed. by JC Beall. Oxford University Press.
- Beaney, Michael, ed. (1997). *The Frege Reader*. Blackwell.
- Blanchette, Patricia A. (2012). *Frege’s Conception of Logic*. Oxford University Press.
- Braun, David and Theodore Sider (2007). “Vague, So Untrue”. In: *Noûs* 41.2, pp. 133–156.
- Broome, John and Wlodek Rabinowicz (1999). “Backwards Induction in the Centipede Game”. In: *Analysis* 59.264, pp. 237–242.
- Collins, John and Achille C. Varzi (2002). “Unsharpenable vagueness”. In: *Philosophical Topics* 28.1, pp. 1–10.
- Dummett, Michael (1991). *Frege: Philosophy of Mathematics*. 171. Harvard University Press, p. 223.
- Fine, Kit (1975). “Vagueness, Truth and Logic”. In: *Synthese* 30.3/4, pp. 265–300.
- Fodor, Jerry A. and Ernest Lepore (1996). “What Cannot Be Evaluated Cannot Be Evaluated and It Cannot Be Supervalued Either”. In: *Journal of Philosophy* 93.10, pp. 516–535.
- Frege, Gottlob (1879). “Begriffsschrift”. In: *The Frege Reader*. Ed. and trans. by Michael Beaney. Blackwell, 1997, pp. 47–78.
- (1892). “On Sinn and Bedeutung”. In: *The Frege Reader*. Ed. by Michael Beaney. Trans. by Max Black. Blackwell, 1997, pp. 151–171.

## BIBLIOGRAPHY

---

- Frege, Gottlob (1903). “Grundgesetze der Arithmetik, Volume II”. In: *The Frege Reader*. Ed. by Michael Beaney. Trans. by Peter Geach. With additions trans. by Michael Beaney. Blackwell, 1997, pp. 258–289.
- (1964). “On the Scientific Justification of a Concept-Script”. Trans. by James M. Bartlett. In: *Mind*. New Series 73.290, pp. 155–160.
- (1969). “Comments on Sinn and Bedeutung”. In: *The Frege Reader*. Ed. by Michael Beaney. Trans. by Peter Long and Roger White. Blackwell, 1997, pp. 172–180.
- (1971). “On the Law of Inertia”. In: *Studies in History and Philosophy of Science Part A* 2.3, pp. 195–212.
- (1979). *Posthumous Writings*. Ed. by Hans Hermes, Friedrich Kambartel, and Friedrich Kaulbach. Trans. by P. Long and R.M. White. Oxford: Basil Blackwell.
- Gabriel, Gottfried et al., eds. (1980). *Gottlob Frege: Philosophical and Mathematical Correspondence*. Blackwell.
- Geach, Peter (1980). *Reference and Generality: An Examination of Some Medieval and Modern Theories*. Third edition. Cornell University Press. Chap. 7, pp. 191–218.
- Heck, Richard G. (2011). *Frege’s Theorem*. Oxford University Press. Chap. 5, pp. 111–126.
- Kamp, Hans (1981). “The Paradox of the Heap”. In: *Aspects of Philosophical Logic*. Ed. by U. Mönnich. Dordrecht, pp. 225–277.
- Kemp, Gary (1996). “Frege’s Sharpness Requirement”. In: *Philosophical Quarterly* 46.183, pp. 168–184.
- Kerényi, Karl (1959). *The heroes of the Greeks*. Myth and man. London: Thames and Hudson.
- Kripke, Saul A. (1975). “Outline of a Theory of Truth”. In: *Journal of Philosophy* 72.19, pp. 690–716.
- Lakoff, George (1973). “Hedges: A Study in Meaning Criteria and the Logic of Fuzzy Concepts”. In: *Journal of Philosophical Logic* 2.4, pp. 458–508.
- Lewis, David (1986). *On the Plurality of Worlds*. Blackwell Publishers.
- Linnebo, Øystein (2010). “Pluralities and Sets”. In: *Journal of Philosophy* 107.3, pp. 144–164.

- Morreau, Michael (1999). "Supervaluation Can Leave Truth-Value Gaps After All". In: *Journal of Philosophy* 96.3, pp. 148–156.
- Priest, Graham (2003). "A site for sorites". In: *Liars and Heaps*. Ed. by JC Beall. Oxford University Press, pp. 9–23.
- Puryear, Stephen (2013). "Frege on Vagueness and Ordinary Language". In: *The Philosophical Quarterly* 63.250, pp. 120–140.
- Quine, W. V. (1962). "Paradox". In: *Scientific American* 206, pp. 84–96.
- Raffman, Diana (1994). "Vagueness Without Paradox". In: *Philosophical Review* 103.1, pp. 41–74.
- Sainsbury, Mark (1996). "Concepts Without Boundaries. Inaugural lecture delivered at King's College London, 1990." In: *Vagueness: A Reader*. Ed. by Rosanna Keefe and Peter Smith. MIT press.
- (2013). "Lessons for Vagueness From Scrambled Sorites". In: *Metaphysica* 14 (2), pp. 225–237.
- Shapiro, Stewart (2008). "Reasoning with Slippery Predicates". In: *Studia Logica* 90.3, pp. 313–336.
- Unger, Peter (1979). "There Are No Ordinary Things". In: *Synthese* 41.2, pp. 117–154.
- Varzi, Achille C. (2007). "Supervaluationism and Its Logics". In: *Mind* 116.463, pp. 633–675.
- Weiner, Joan (2007). "What's in a Numeral? Frege's Answer". In: *Mind* 116.463, pp. 677–716.
- Williamson, Timothy (1994). *Vagueness*. London: Routledge.
- (2000). *Knowledge and its Limits*. Oxford University Press. Chap. 5, pp. 114–134.
- Williamson, Timothy and Peter Simons (1992). "Vagueness and Ignorance". In: *Proceedings of the Aristotelian Society, Supplementary Volumes* 66, pp. 145–177.
- Wittgenstein, Ludwig (1974). *Philosophical Grammar*. Blackwell.
- Wright, Crispin (1975). "On the Coherence of Vague Predicates". In: *Synthese* 30.3-4, pp. 325–65.

## BIBLIOGRAPHY

---

- Wright, Crispin (1976). "Language-Mastery and the Sorites Paradox". In: *Truth and Meaning: Essays in Semantics*. Ed. by Gareth Evans and John Henry McDowell. Clarendon Press.
- (1995). "The Epistemic Conception of Vagueness". In: *Southern Journal of Philosophy* 33.S1, pp. 133–160.
- Yablo, Stephen (2004). "Circularity and Paradox". In: *Self-Reference*. Ed. by Vincent F. Hendricks & Stig Andur Pedersen Thomas Bolander. Csi Publications, pp. 139–157.
- Zalta, Edward N. (2013). "Frege's Theorem and Foundations for Arithmetic". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Fall 2013.