

Introduction

Ordinarily we say that *persons* are vague if they regularly misplace their car keys, stare blankly into the fridge having forgotten why they opened it, and so on. In philosophy, however, the term ‘vague’ has a different use. It applies primarily (although not exclusively) to *predicates*.¹ Amongst all predicates, the vague ones are usually singled out in one or more of three ways:

Borderline cases. There are some persons to whom the predicate ‘is tall’ clearly applies (e.g. most professional basketball players) and some to whom it clearly does not apply (e.g. most professional jockeys), but then there are other persons to whom it is unclear whether or not the predicate applies (I’m sure you know some of them). When asked whether such a person is tall, we tend to react with some sort of *hedging response*: “sort of”; a shrug and a certain sort of scowl or exhalation of breath; a blank look; etc. Call these persons *borderline cases* for ‘tall’. Other predicates—for example ‘is a prime number’—do not have borderline cases. Julius Caesar is not a prime number, seven is a prime number, eight is not a prime number, and so on: there is nothing at all of which it is unclear whether it is a prime number.² One standard way of drawing the distinction between vague and non-vague (or precise) predicates is to say that vague predicates have borderline cases, while precise predicates do not.³

Blurred boundaries. Imagine a line drawn around all the things to which a given predicate applies. One typical characteristic of vague predicates is that we get a blurry, ill-defined line. For example, there is no sharp boundary between the things to which ‘is heavy’ applies and the things to which it does not apply. For a precise predicate, by contrast, we get a sharp line, cleanly separating the things to which the predicate applies from the rest.⁴

¹Predicates are items of language—e.g. ‘is tall’, ‘is happy’, ‘is running’—which go together with names—e.g. ‘Bill’, ‘Ben’, ‘Alice’—or definite descriptions—e.g. ‘the tallest woman in the room’, ‘the inventor of post-it notes’, ‘the man who threw the egg’—to form sentences—e.g. ‘Bill is tall’, ‘The inventor of post-it notes is happy’, ‘Alice is running’.

²For large numbers, we may suppose that we can use a computer to help us determine the answer.

³The borderline case characterisation can be traced at least as far as Peirce [1902].

⁴The blurred boundaries characterisation can be traced at least as far as Frege’s statement that if we represent concepts in extension by areas on a plane, then vague concepts do not have sharp

Sorites paradoxes. Consider a hirsute person. He is not bald. Were we to remove a single one of his hairs, he still would not be bald. So too if we removed another, and so on—because *one hair* cannot ever make the difference between baldness and lack thereof. But our man has only a finite number of hairs, so eventually he will have none left. By our reasoning, he will still not be bald—whereas of course a person with no hairs on his head *is* bald. So something has gone wrong somewhere. This is an instance of the *Sorites paradox*. Note that the problem does not arise for all predicates. Instead of ‘is bald’, take the predicate ‘has ten or more hairs on his head’. Consider a hirsute person. He has ten or more hairs on his head. Were we to remove a single one of his hairs, he would still have ten or more hairs on his head. So too if we removed another, and another—but *not* ‘and so on and on’, because one hair *can* make the difference between having ten or more hairs, and having nine or less. In this case we get no paradox. One standard way of drawing the distinction between vague and precise predicates is to say that vague predicates generate Sorites paradoxes, while precise predicates do not. That is, where ‘is *P*’ is vague, we will typically be able to imagine a so-called Sorites series of objects, ranging from one which is clearly *P* to one which is clearly not *P*, but where we *also* feel strongly inclined to say that for any object in the series, if it is *P* then so are its neighbours.⁵

So we are interested in predicates which admit of borderline cases, which draw blurred boundaries, and which generate Sorites paradoxes. *Why* are we interested in such predicates? For at least two reasons. First, vagueness is *ubiquitous*. Most of the predicates in our language are vague. Indeed, it is very hard to think of a predicate which is not at all vague. Some predicates have a wider class of borderline cases than others, but it is quite a challenge to think of predicates outside mathematics and the hard sciences which admit of no possibility of borderline cases—whose boundaries are not in the slightest bit blurry. (Try it now!) So vagueness is something we cannot avoid. If we want to understand our language and our world, we need to understand vagueness. Second, there is immense potential for many *other* research programmes to be significantly advanced by progress on vagueness. When a key concept in a certain area is vague—and thanks to the ubiquity of vagueness, key concepts in many areas *are* vague—understanding how its vagueness works can have crucial payoffs.

So we know which predicates we want to study, and why we want to study them. What exactly do we hope to gain from the study? What is our goal? What we ultimately want to gain from the study of vagueness is an account of the semantics of vague predicates—an account of what the meaning of such a predi-

boundaries, but rather fade off into the background [*Grundgesetze* vol. II, §56; translated in Beaney 1997, 259].

⁵More precise characterisations of Sorites series and Sorites paradoxes will be presented below in §3.5.4.

cate consists in, and of how (if at all) it differs from the meaning of a precise predicate. The development of modern logic (also known as classical logic) and model theory in the 19th and 20th centuries led to immense progress in our understanding of language. However, the pioneers of these developments—notably Gottlob Frege and Alfred Tarski—were primarily concerned with *mathematical* language, and they ignored vagueness. Frege in particular set vagueness to one side as not worthy—or not susceptible—of treatment. He considered vagueness a defect of natural language, to be banished from a logically perfect scientific language. The *problem* with this attitude is that mathematical language is a special case: it is completely precise. Outside mathematics, virtually all of our language is vague, to a greater or lesser extent. Beginning in real earnest in the second half of the 20th century, this realisation that vagueness is ubiquitous in natural language led to a desire to approach ordinary language in the same spirit in which Frege approached mathematical language—in a bid not to eliminate vagueness, but to understand and adequately represent it. This led to the development of a large number of non-classical logics, and associated theories of vagueness (where by a *theory* of vagueness I mean a formal logical core, together with a package of motivating and explanatory philosophical views). There are now many non-classical logics, and associated theories of vagueness, on the market. So in one sense we have advanced greatly since Frege’s time: we now have many systems of semantics designed specifically to accommodate vague language. In another sense, however, we are still no closer to our goal of obtaining an account of what the meaning of a vague predicate consists in—for we have *many* accounts, and no apparent way of deciding which (if any) of them is *correct*.

Part I: Foundations

As a first step towards rectifying this situation, in Chapter 2—building on basic preliminary material which it is the purpose of Chapter 1 to present—I give an overview of the space of possible theories of vagueness, and show where existing theories live in this space. For we cannot begin to select a theory of vagueness until we fully understand how the different theories differ from one another: not just over particular matters of detail, but at a more fundamental and illuminating level of analysis. Accordingly, Chapter 2 provides a conceptual map of theories of vagueness, setting out the important axes along which different theories vary, and then giving the coordinates of existing and future theories on these axes. At the centre of the map is the classical view of vagueness, according to which vague language is—from a semantic point of view—exactly the same as precise language. This view applies the standard model theory for precise mathematical language to vague language. I begin by isolating the key elements of this classical model theory: these key elements give the axes of the space. With the classical view at

the origin, other theories are located according to *which* of the key elements they modify, and then, at a finer level of analysis, according to *how* they modify them.

It is in this chapter that a distinction which plays a central role in this book first emerges: the distinction between worldly vagueness and semantic indeterminacy. Here is a quick sketch of how it emerges. Model theory involves two aspects: a part which represents a *language*, and a part—a *model*—which represents the *world*. The process whereby items of language gain meaning is represented by matching up a language with a model—which is called giving an *interpretation* of the language. The underlying thought is that our words gain meaning when we use them to talk about the world. But language is versatile, and words could mean many things. Consider the sentence ‘The house is on the hill’. The symbol ‘house’ *could* have been used to mean what the symbol ‘cat’ actually means, and the symbol ‘hill’ could have been used to mean what the symbol ‘mat’ actually means, so there is an interpretation of the sentence on which it means that the cat is on the mat. But the interpretation on which the sentence means that the house is on the hill is privileged: it corresponds to what the sentence actually means. Such an interpretation is called the *intended* interpretation. In isolating the key elements of the classical model theory for precise language, I make a broad distinction between the *internal* nature of classical models, and the *external* part of the classical story, according to which each discourse has a *unique* intended interpretation. The classical picture tells us on the one hand that the world is of such and such a sort (it has a structure which can be represented by a classical model, and classical models have such and such internal features), and on the other hand that the meaning of anything we say is correctly specified by giving a *single* such model. This distinction yields two broad ways in which a theory of vagueness can differ from the classical view: it can replace classical models with non-classical models with different internal features, or it can alter the external story according to which only one model comes into play in describing what some utterance actually means. This distinction between two ways of departing from classical model theory underlies the distinction between worldly vagueness and semantic indeterminacy. Vagueness and indeterminacy are ruled out of the classical picture in *two* places. First, there is no vagueness in the world. Classical models are entirely precise: they represent the world as a crisp set of objects and properties such that for each property and each object, the object either definitely possesses the property, or definitely does not possess it. Second, there is no semantic indeterminacy—no vagueness in the relationship between language and the world. Each sentence has a *unique* intended model: so there is no vagueness about what any sentence means. A view which differs from the classical picture in the internal way (i.e. replaces classical models with non-classical models with different internal features) opens up the possibility of *worldly vagueness*, while a view which differs in the external way (i.e. denies that only one model comes into play in describing what

some utterance actually means) opens up the possibility of *semantic indeterminacy*. For example, the view of vagueness built on fuzzy set theory (see §2.2.1) countenances worldly vagueness. It tells us that when we utter a vague predicate, we pick out a unique property in the world—the extension of our predicate on the unique intended interpretation of our utterance—but this property is *inherently* vague, in the sense that some objects possess it, some do not possess it, and others possess it to various intermediate degrees. This is a matter of how things are out there in the world (on this view)—it is not a matter of the relationship between language and the world. In contrast, semantic indeterminacy enters when we have a view which denies that an utterance has a unique intended interpretation. On this sort of view, when we utter a vague predicate, there are *several* properties we might be picking out—the extensions of this predicate on the various equally intended (or equally not-unintended) interpretations—and there is no fact of the matter as to which, in particular, we mean. Here we have indeterminacy in the relationship between language and the world.⁶

Some surprising and significant results emerge from the survey of theories of vagueness. One is that two quite different views have been conflated in the literature under the heading ‘supervaluationism’. I call these two views ‘supervaluationism’ and ‘plurivaluationism’. This distinction is particularly important because ‘supervaluationism’ is widely regarded as the front runner amongst existing theories of vagueness—but once we clearly distinguish the two views that have been run together under this heading, we will be in a much better position to discern the real disadvantages of each.

Part II: Vagueness

With a clear overview of theories of vagueness in hand at the end of Chapter 2, we still face the question of deciding which theory is correct. The literature abounds with objections to particular theories, and these have led to the outright rejection of a number of views—but still, several broad *types* of theory remain viable (for example epistemicism and degree-theoretic approaches, to name just two). Within each type, there is general agreement that some particular versions are better than others—but there are no widely accepted criteria which can decide between the types themselves. Theorists of vagueness have thus found themselves divided into camps. This is a very unsatisfactory situation. Given that the different

⁶In order to draw substantive conclusions about the location of vagueness—in the world itself, or in the relationship between language and the world—from a system of model theory, we must take a *literal* attitude towards model theory. That is, we must regard a model theory for (part of) a language as giving a literal (although not necessarily complete) description of the relationship between that language and the world. I present reasons for taking such a literal attitude towards model theories for vague language (in the context of a study such as this one) in §2.1.3.1.

types of view present completely different pictures of the relationship between vague language and the world, they cannot all be correct. We need some way of deciding which type of theory is right. We need some criteria for determining what the *general form* of the correct theory of vagueness should be.

There is another loose end that also cries out to be tied up at this point. We have three informal characterisations of vague predicates: they give rise to borderline cases, their extensions have blurry boundaries, and they generate Sorites paradoxes. These three characterisations seem to be quite closely related to one another—yet not so closely related that they are merely three ways of saying essentially the same thing. So what is the relationship between them? Rather than three piecemeal characterisations of vagueness, it would be desirable to have a proper definition of vagueness: a crisp statement of what is of the essence of vagueness—of what vagueness ultimately consists in—given which, we can see *why* vague predicates have borderline cases, generate Sorites paradoxes, and draw blurred boundaries.

My strategy is to link these two tasks: the task of finding the correct theory of vagueness, and the task of finding an adequate definition of vagueness. Chapter 3 is concerned with the latter task. In that chapter I discuss what we should expect from a definition of vagueness, critically discuss existing definitions, and then present a definition of vagueness according to which—roughly—a predicate F is vague just in case for any objects a and b , if a and b are very similar in respects relevant to the application of F , then the sentences Fa and Fb are very similar in respect of truth. So for example, ‘is tall’ will be vague just in case for any persons a and b , if a and b are very similar in height, then the sentences ‘ a is tall’ and ‘ b is tall’ are very similar in respect of truth. The import of this definition can best be grasped by comparing it with the claim that vague predicates are *tolerant*. A predicate F is tolerant with respect to ϕ if there is some positive degree of change in respect of ϕ that things may undergo, which is “insufficient ever to affect the justice with which F is applied to a particular case” [Wright, 1975, 334]. So for example, ‘is tall’ will be tolerant with respect to height just in case for any persons a and b , if a and b are very similar in height, then there is no difference in the applicability of the predicate ‘is tall’ to a and b —that is, the sentences ‘ a is tall’ and ‘ b is tall’ are *exactly the same* in respect of truth. The great problem with the claim that any predicate F is tolerant is that, when conjoined with the claim that we can construct a Sorites series for the predicate F , it leads to contradiction—in particular, to the claim that each object in the Sorites series both is and is not F . My definition of vagueness is a weakening of the claim that vague predicates are tolerant: if a and b are very similar in height, then there need not be *no* difference in the applicability of the predicate ‘is tall’ to a and b —that is, the sentences ‘ a is tall’ and ‘ b is tall’ need not be *exactly the same* in respect of truth—but there cannot be *much* difference in the applicability of the

predicate 'is tall' to *a* and *b*, and so the sentences '*a* is tall' and '*b* is tall' must be *very similar* in respect of truth. In the remainder of Chapter 3 I discuss the question of extending this definition to cover vagueness of many-place predicates, of properties and relations, and of objects, and I then explore the advantages of this definition, some of the most important of which are that it captures the intuitions which motivate the thought that vague predicates are tolerant, *without* leading to contradiction, and that it yields a clear understanding of the relationships between Sorites-susceptibility, blurred boundaries and borderline cases.

In Chapter 4 I turn to the task of determining what type of theory of vagueness we need. Having clearly presented the different types of theory in Chapter 2, and having presented a crisp definition of vagueness in Chapter 3, my strategy is to ask which types of theory can accommodate vague predicates—that is, predicates which satisfy the definition given in Chapter 3. It is here, then, that the link is made between the task of finding an adequate definition of vagueness, and the task of finding the correct theory of vagueness. When vagueness is characterised informally in terms of borderline cases, blurred boundaries, and Sorites-susceptibility, *all* the main existing types of theory of vagueness can be seen as accommodating vagueness—simply because the informal characterisations are so loose. This leads to the 'too many theories' problem: the existing theories cannot all be right, as they conflict with one another; yet how to choose between them, if they all accommodate vagueness perfectly well? The situation changes, however, once we have to hand a sharp definition of the core property underlying the various surface phenomena standardly used to characterise vagueness. When we now ask whether each type of theory allows for the existence of predicates possessing the feature isolated in Chapter 3 as being of the essence of vagueness, it turns out that the answer is No. Only one type of theory does: the type which countenances *degrees of truth*.

The basic idea of degrees of truth is that while some sentences are true and some are false, others possess intermediate truth values: they are truer than the false sentences, but not as true as the true ones. So, for example, if we line up a series of persons, ranging from one who is 7 feet in height to one who is 4 feet in height, in increments of a fraction of an inch, and then move along the series saying of each person in turn, 'This person is tall', the idea is that our statements start out quite true, then gradually get less and less true, until they end up quite false.

Thus, from the project of defining vagueness, an answer emerges to the question of what the general form of the correct theory of vagueness must be: it must be one which countenances degrees of truth.⁷ In order to reach the overall conclu-

⁷Clearly, in order for the *definition* of vagueness to play the role of deciding between *theories* of vagueness, we must reject the widespread idea that a definition should be theory-neutral—that

sion of Chapter 4—that we need a theory of vagueness that countenances degrees of truth—I need to do several things in the course of the chapter. First, I show that of the types of theory of vagueness distinguished in Chapter 2, only those which countenance degrees of truth can accommodate predicates which satisfy the definition of vagueness proposed and defended in Chapter 3. Second, I consider and reject two different strategies which non-degree theorists might employ to avoid my conclusion. The first strategy is to propose an *error theory* of vagueness. Non-degree theorists might *agree* with the definition of Chapter 3, but still maintain the correctness of their semantic theory. The resulting position would be as follows: “For a predicate to be vague, it has to satisfy the definition of Chapter 3; my semantic theory does not allow for the existence of predicates that satisfy the definition of Chapter 3; but my semantic theory is correct; therefore there are no vague predicates.” I offer reasons for rejecting such an error theory of vagueness. The second strategy is to *reject* the definition of Chapter 3, to propose in its place an alternative definition which *is* compatible with a given non-degree theory, and to argue that this alternative definition has all the advantages which I claim for my definition in Chapter 3. I consider, and give reasons for rejecting, several proposals along these lines, including the proposal that a predicate F is vague just in case for any objects a and b , if a and b are very similar in respects relevant to the application of F , then the sentences Fa and Fb are very similar in respect of *assertibility* (rather than *truth*).

Part III: Degrees of Truth

The upshot at this point is that we need a theory which countenances degrees of truth. This tells us the *type* of theory we need, but does not lead us to a particular theory. The best-known degree theory is the one based on fuzzy logic and set theory. It replaces the two classical truth values True and False with infinitely many degrees of truth, which may be thought of as percentage values: thus a sentence might be only 45% true, or 98% true, and so on. In Chapters 5 and 6 I consider objections to the fuzzy view of vagueness in particular, and to degree-theoretic treatments of vagueness in general. In some cases, I argue that the objections do not carry weight. In other cases, I propose modifications and/or additions to the standard fuzzy view, in order to overcome the objections.

The main objections covered in Chapter 5 are as follows. The very idea of truth coming in degrees is in some way a mistake! I show that it is not. The fuzzy theory involves an objectionable violation of classical logic! In response to this objection, I propose a new account of logical consequence in the fuzzy setting that allows

is, that the definition of P must be something on which all the candidate theories of P can agree. I discuss this issue in Chapter 3.

us to derive a *classical* consequence relation from fuzzy semantics. That is, while the *semantics* for vagueness that I propose is non-classical, it validates classical *logic*. Degrees of truth cannot be integrated with key developments elsewhere in philosophy of language, outside the study of vagueness! I show that this is not correct. As part of my response to this objection, I propose a detailed new account of the relationship between degrees of truth and degrees of belief. Degree theories, such as the fuzzy theory, which treat the logical connectives as truth functions, cannot account for ordinary usage of sentences about borderline cases! I show that this objection is mistaken. Finally, denying bivalence—the view that every sentence is true or false, with no middle way (or ways, in the case of degrees of truth)—leads to contradiction! I show that it does not.

In Chapter 6 I turn to what is arguably the major objection to the fuzzy view, viz., that it “imposes artificial precision. . . though one is not obliged to require that a predicate either definitely applies or definitely does not apply, one *is* obliged to require that a predicate definitely applies to such-and-such, rather than to such-and-such other, degree (e.g. that a man 5 ft 10 in tall belongs to *tall* to degree 0.6 rather than 0.5)” [Haack, 1979, 443]. In response to this objection, I propose a new view—which I call *fuzzy plurivaluationism*—which departs from the standard fuzzy view by denying that each discourse has a unique intended (fuzzy) interpretation. This view thus incorporates *both* worldly vagueness (in the fuzzy models) *and* semantic indeterminacy (in the lack of a unique intended interpretation). The overall conclusion of the book is that fuzzy plurivaluationism is the correct theory of vagueness.

Conclusion

We have covered a lot of ground, and I will not recapitulate every step. It might, however, be useful to conclude by retracing one key line of thought, and showing how it leads to the positive view that I have presented and defended: fuzzy plurivaluationism.

The line of thought begins with the observation that there are *two* problems with the epistemicist view of vagueness—specifically, with its positing of a particular change point in a Sorites series for P at which the claim ‘this object is P ’ goes from being true to being false. The first problem is what I call the *jolt problem*. The problem is the *nature* of the semantic shift posited by the epistemicist, and in particular the dramatic difference between the semantic statuses of the claims ‘this object is P ’ (which is true) and ‘that object is P ’ (which is false), where these claims concern the objects on either side of the change point—objects which are very similar in all respects relevant to the application of the predicate P . In short, the epistemicist thinks that as we go down the Sorites series saying ‘this is P ’ of each object in turn, there is at some point a sudden jolt, as our claims crash all the way from true to false in one step—and this has been thought to be implausible. The second problem is what I call the *location problem*. The problem concerns the *fixing of the location* of the change point, and in particular the fact that we cannot see how our own usage of language (together with facts about our causal connections to our environment, referential eligibility, simplicity and so on) could fix it to be at any particular point in the series. The epistemicist thinks that there is a number n such that the n th object in the series is the last one to which P truly applies—but why, it has been objected, not $n - 1$ or $n + 1$? What is it about our usage of P that gives it a meaning which singles out a unique n in this way?

In the literature, these two problems have tended to be run together under the heading ‘higher-order vagueness’. For example, when someone wants to object to (say) a truth-gap view of vagueness, on the grounds that, while it avoids a point in the Sorites series at which the claim ‘this is P ’ crashes from true to false, it does not avoid dramatic semantic shifts altogether—for it posits a dramatic difference between the semantic statuses of the claims Pa and Pb , where a and b straddle the boundary between the positive (negative) cases and the borderline cases, in that Pa is true (false) while Pb lacks a truth value, even though a and b are very similar

in all respects relevant to the application of P —she typically does so by saying that the gappy view falls foul of the problem of higher-order vagueness. And when, for example, someone wants to object to (say) a supervaluationist view, on the grounds that while it avoids positing a *unique* number n such that the n th object in the series is the last one to which P applies, it does not avoid problems of meaning-determination altogether—for it posits a precise *set* of admissible precisifications of vague language, when it is extremely hard to see how our practice could determine a unique such set—he typically does so by saying that the supervaluationist view falls foul of the problem of higher-order vagueness. But the two objections should not be run together under one name, for the first is an instance of the jolt problem, and the second is an instance of the location problem—and these are fundamentally different problems. The jolt problem is intimately connected to vagueness. I have argued that it is of the essence of vagueness that dramatic semantic shifts do not occur in Sorites series: for if P is vague, and a and b are very similar in P -relevant respects (and adjacent items in a Sorites series for P always are very similar in P -relevant respects) then Pa and Pb are very similar in respect of truth. The location problem, on the other hand, is *not* intimately connected to vagueness. I have argued that it is a manifestation of more general worries about what fixes the meaning of our language—worries which manifest in areas having nothing to do with vagueness, such as Quine’s problem of the indeterminacy of translation and Kripkenstein’s sceptical problem.

So we have two quite different problems to solve—and this is the origin of the two key features of fuzzy plurivaluationism. First, its positing of degrees of truth is a response to the jolt problem. I argued that no theory can allow for the existence of predicates which both have associated Sorites series and satisfy Closeness—i.e. no theory can avoid the jolt problem—unless it countenances degrees of truth. Second, its positing of semantic indeterminacy—of a lack of a unique intended interpretation of vague discourse—is a response to the location problem. I argued that the *vagueness* of a discourse would not be impugned simply because it had a unique intended fuzzy interpretation—but that as a matter of fact, it seems that our meaning-fixing practices do not in general suffice to determine a unique intended interpretation of vague language, and so we must countenance semantic indeterminacy. Other theories of vagueness can solve *one* of these two problems. For example, classical plurivaluationism can solve the location problem, and the original fuzzy theory, which posits degrees of truth without semantic indeterminacy, can solve the jolt problem. However only fuzzy plurivaluationism—which posits both degrees of truth and semantic indeterminacy—can solve both problems.