Forthcoming in M. Loux and D. Zimmerman, eds., *The Oxford Handbook of Metaphysics*, Oxford: Oxford University Press

Vagueness in Reality

Timothy Williamson

1. Introduction

When I take off my glasses, the world looks blurred. When I put them back on, it looks sharp-edged. I do not think that the world really was blurred; I know that what changed was my relation to the distant physical objects ahead, not those objects themselves. I am more inclined to believe that the world really is and was sharp-edged. Is that belief any more reasonable than the belief that the world really is and was blurred? I see more accurately with my glasses on than off, so visual appearances when they are on have some cognitive priority over visual appearances when they are off. If I must choose which kind of visual appearance to take at face value, I will choose the sharp-edged look. But what should I think when I see a mist, which looks very blurred however well I am seeing? Indeed, why choose to take any of the looks at face value? Why not regard all the choices as illegitimate projections of ways of seeing the world onto the world itself?

Such questions arise for thought and language as well as for perception. They concern vagueness, susceptibility to borderline cases in which a judgement is neither clearly correct nor clearly incorrect. For example, Mount Everest has vague boundaries: some rocks are neither clearly part of Everest nor clearly not part of Everest. Is Everest therefore a vague object? Or is only the name 'Everest' vague? If the name is vague, is it a vague object, since names are objects too? In what sense, if any, is all vagueness mind-dependent?

Raised in a theoretical vacuum, such questions quickly produce confusion. We make more progress by teasing out the metaphysical consequences of vagueness within a systematic framework for reasoning with vague concepts. Sections 2 and 3 explain two main proposals for accommodating vagueness by modifying the traditional dichotomy between truth and falsity and consequently rethinking logic: fuzzy logic and supervaluationism. Within these frameworks, we wish to ask 'Is reality vague?'. What does that mean? Section 4 uses the idea of a state of affairs to clarify the question. Sections 5 and 6 show how supervaluationism and fuzzy logic embody opposite answers to it. Of course, a semantic theory alone cannot answer the question, because it is not a question about language, but it can say what reality must be like for our claims about vagueness to be true. Section 7 considers whether objects, properties and relations can be vague. Section 8 concerns the special case of vague identity. Section 9 introduces an epistemic theory of vagueness as an alternative perspective on all these questions.

This chapter is not intended as a general introduction to the philosophy of vagueness, much of which falls within the philosophy of language and epistemology rather than metaphysics.[1] Instead, the aim is to elucidate the metaphysical significance of some subtle issues in the logic of vagueness.

2. Fuzzy semantics[2]

A common proposal is to adjust logic to vagueness by smoothing out the classical dichotomy of truth and falsity into a continuum of degrees of truth. On this view, as the sun sets, 'It is dark' starts off definitely false, gradually increases in degree of truth, and ends up definitely true. For convenience, degrees of truth are usually identified with the continuum of real numbers between 0 and 1 inclusive: 0 is definite falsity, 1 is definite truth. When 'It is dark' has degree of truth 0.5, it is just as true as false. This is the basis of fuzzy logic. What distinguishes the fuzzy approach is not merely that it postulates degrees of truth, for proponents of other approaches such as supervaluationism can postulate them too (see section 3). Rather, fuzzy semantics is distinctive because it gives degrees of truth the same structural role as truth and falsity play in standard truth-conditional semantics. In particular, just as standard semantics gives the meaning of a logical connective by stating how the truth-value of a complex sentence composed from simpler sentences with that connective depends on the truth-values of those simpler sentences, so fuzzy semantics gives the meaning of the connective by stating how the degree of truth of the complex sentence depends on the degrees of truth of the simpler sentences.

Since we want to express generalizations as well as particular claims, we shall need a language with variables as well as constants. As in classical semantics, variables refer not absolutely but relative to assignments of appropriate values.[3] Call the referent of an expression e relative to an assignment a of values to all variables '$\mathrm{Ref}_a(e)$'. If e is a variable, $\mathrm{Ref}_a(e)$ is simply $a(e)$, the value a assigns to e. If e is a constant, $\mathrm{Ref}_a(e)$ is independent of a. Call the degree of truth of a formula α relative to an assignment a '$\mathrm{Val}_a(\alpha)$'.

In an atomic sentence $Ft_1...t_n$, F is an n-place predicate and $t_1, ..., t_n$ are singular terms. We call the referents of singular terms 'objects' without prejudice to their nature. In two-valued semantics, an n-place predicate maps each n-tuple of objects to a truth-value. Thus the 2-place predicate 'kisses' maps the ordered pair <John, Mary> to truth if John kisses Mary and to falsity otherwise. Likewise in fuzzy semantics, an n-place predicate maps each n-tuple of objects to a degree of truth; for example, 'kisses' might map <John, Mary> to 0.5 if Mary draws back before it is clear whether John has succeeded in kissing her. Thus the degree of truth of the atomic sentence relative to an assignment a is the value to which the function $\text{Ref}_a(F)$ maps the n-tuple of objects $<\text{Ref}_a(t_1),...,\text{Ref}_a(t_n)>$. Formally:

FUZZYatom $\quad \text{Val}_a(Ft_1...t_n) = \text{Ref}_a(F)(\text{Ref}_a(t_1),...,\text{Ref}_a(t_n))$

Classical semantics treats a logical operator such as negation (~), conjunction (&) or disjunction (∨) as truth-functional: the truth-value of a complex sentence consisting of such an operator applied to one or more simpler sentences is a function of the truth-values of those simpler sentences, displayed by the truth-table for that operator. Similarly, fuzzy semantics treats those operators as degree-functional: the degree of truth of the complex sentence is a function of the degrees of truth of the simpler sentences. For negation the obvious proposal is:

FUZZY~ $\quad \text{Val}_a(\sim\alpha) = 1\text{-Val}_a(\alpha)$

If α is definitely false, it has degree of truth 0, so ~α has degree of truth 1 and is definitely true. If α is definitely true, it has degree of truth 1, so ~α has degree of truth 0 and is definitely false. Thus, if we equate truth and falsity with degree of truth 1 and 0 respectively, the classical truth-table for ~ emerges as a special case of FUZZY~. But as 'It is dark' gradually increases in degree of truth, so 'It is not dark' gradually decreases. At the halfway point, both statements have degree of truth 0.5; they characterize the situation equally well. The case is perfectly borderline.

So far, degrees of truth look formally just like probabilities (the probability of ~α is one minus the probability of α). Not so for other operators. The probability of a conjunction is not determined by the probabilities of its conjuncts. For if both conjuncts have probability 0.5, the probability of the conjunction may be anywhere between 0 and 0.5. For a fair coin on a given toss, the conjunction 'The coin will come up heads and the coin will come up tails' has probability 0 while both its conjuncts have probability 0.5; the repetitive conjunction 'The coin will come up heads and the coin will come up heads' has probability 0.5 while both its conjuncts have probability 0.5 too. But fuzzy semantics calculates the degree of truth of a conjunction from the degrees of truth of its conjuncts. The standard proposal is that it is their minimum:

FUZZY& $\mathrm{Val}_a(\alpha \,\&\, \beta) = \min\{\mathrm{Val}_a(\alpha), \mathrm{Val}_a(\beta)\}$

The classical truth-table for & emerges as a special case of FUZZY&, for it implies that the conjunction has degree of truth 1 if both conjuncts have degree of truth 1 and degree of truth 0 if at least one conjunct has degree of truth 0. FUZZY& has intuitively attractive features. It implies that adding a conjunct never increases the degree of truth of a conjunction (the expanded conjunction entails the original one), that increasing the degree of truth of a conjunct never decreases the degree of truth of a conjunction, and that repeating a conjunct makes no difference to the degree of truth of a conjunction. One can easily show that FUZZY& gives the *only* function from the degrees of truth of the conjuncts to the degree of truth of the conjunction that satisfies those desiderata.

Nevertheless, FUZZY& has counterintuitive consequences. Suppose, for example, that the twins Jack and Mack are balding in the same way. Their scalps are in exactly the same state; they are bald to exactly the same degree. However far the process has gone, the claim 'Jack is bald and Mack isn't' (*Bj* & ~*Bm*) is not perfectly balanced between truth and falsity; intuitively, it is false, or at least much closer to falsity than to truth. The conjunction is not merely conversationally misleading, by not speaking symmetrically of the twins, for the conditional 'If Jack is bald then Mack is bald' does not speak of them symmetrically but is intuitively correct. *Bj* & ~*Bm* should receive a degree of truth less than 0.5 at every point of the synchronized balding processes. At the halfway point, *Bj* has degree of truth 0.5. So has *Bm*, for Jack and Mack are in the same state; so has ~*Bm* by FUZZY~. Therefore FUZZY& assigns *Bj* & ~*Bm* degree of truth 0.5, the wrong result. A similar argument shows that FUZZY& assigns the contradiction *Bj* & ~*Bj* degree of truth 0.5 at the halfway point. Yet, intuitively, the contradiction is closer to falsity than to truth. If fuzzy semanticists attempt to avoid these results by using a different function to compute the degrees of truth of conjunctions, they violate the intuitive desiderata that only FUZZY& satisfies. The problem is not in the choice of function from the degrees of truth of the conjuncts to the degree of truth of the conjunction but in the very idea that there is such a function. This kind of objection has long been familiar under the terminology of 'penumbral connections' (Fine 1975: 269-70, Kamp 1975: 131, Williamson 1994: 135-8). Fuzzy logicians have never found a convincing response. We should remember this dark cloud over fuzzy logic as we continue to discuss the approach.

Just as the classical truth-tables assign a conjunction the lowest truth-value of its conjuncts, so they assign a disjunction the highest truth-value of its disjuncts. Just as fuzzy logicians generalize to the minimum degree of truth for conjunctions, so they generalize to the maximum degree of truth for disjunctions:

FUZZY∨ $Val_a(\alpha \vee \beta) = \max\{Val_a(\alpha), Val_a(\beta)\}$

Conjunction and disjunction raise similar issues. Note that FUZZY∨ makes the law of excluded middle a half-truth. When α has degree of truth 0.5, so has $\sim\alpha$; thus by FUZZY∨ so too has $\alpha \vee \sim\alpha$.

What of the quantifiers? A universal generalization resembles the conjunction of its instances. Since some objects lack names, the instances need not correspond to distinct sentences in the language; the conjunction is not to be understood substitutionally. Rather, $\forall v\alpha$ (where α may contain the variable v) is equivalent to the conjunction of α itself under all assignments to v (while the assignments to other variables are held fixed). That is as in classical semantics. Given FUZZY&, one might therefore expect the degree of truth of a universal generalization to be the minimum of the degrees of truth of its instances. However, there is a technical hitch, for if it has infinitely many instances, there may be no minimum degree of truth. For example, the set {0.99, 0.909, 0.9009, ...} has no least member. Nevertheless, although 0.9 is not a member of that set, it is its greatest lower bound (glb): the greatest number not greater than any member of the set. Every set of degrees of truth has a unique greatest lower bound. Accordingly, the proper generalization of FUZZY& is:

FUZZY∀ $Val_a(\forall v\alpha) = \text{glb}\{Val_{a*}(\alpha)$: $a*$ an assignment differing from a at most on $v\}$

Similarly, we extend FUZZY∨ to a rule for the existential quantifier using the analogy between existential generalization and disjunction. As greatest lower bounds replace minima, so least upper bounds (lubs) replace maxima:

FUZZY∃ $\text{Val}_a(\exists v\alpha) = \text{lub}\{\text{Val}_{a*}(\alpha)$: a^* an assignment differing from a at most on $v\}$

So far, we cannot talk *about* vagueness in the formal language. We might want to say that Jack is a borderline case of baldness, neither definitely bald nor definitely not bald. To do so, we introduce the operator Δ, 'definitely'. $\sim\Delta Bj$ & $\sim\Delta{\sim}Bj$ says roughly that Jack is neither definitely bald nor definitely not bald. Δ means 'it is definite that' rather than 'it is definite whether', so $\Delta\alpha$ is false if α is false. The obvious rule is that $\Delta\alpha$ is definitely true when α is definitely true and definitely false otherwise. Since definite truth is degree of truth 1, that amounts to:

FUZZYΔ $\text{Val}_a(\Delta\alpha) = 1$ if $\text{Val}_a(\alpha) = 1$

$\text{Val}_a(\Delta\alpha) = 0$ otherwise

Fuzzy logicians generally use classical metalogic when reasoning in the metalanguage. Thus they accept the metalinguistic instance of the law of excluded middle that either $\text{Val}_a(\alpha) = 1$ or $\text{Val}_a(\alpha) \neq 1$. Consequently, either $\text{Val}_a(\Delta\alpha) = 1$ or $\text{Val}_a(\Delta\alpha) = 0$. Therefore, $\Delta\Delta\alpha \vee \Delta{\sim}\Delta\alpha$ always has degree of truth 1, so $\Delta\alpha$ is not itself vague. Thus nobody is a borderline case of definite baldness. That is counterintuitive. As Jack gradually goes bald, it is no clearer when he becomes definitely bald than it is when he becomes bald. It is sometimes unclear whether he is bald; it is sometimes unclear whether he is definitely bald. This is the problem of higher-order vagueness. If first-order borderline cases should be assigned intermediate degrees of truth, so should second-order borderline cases. No modification of FUZZYΔ by itself would help, for within fuzzy semantics FUZZYΔ defines a perfectly good operator that intuitively is vague. To redefine Δ is to change the subject, not solve the problem. One major unmet challenge facing fuzzy logic is to give an adequate treatment of higher-order vagueness. Presumably, that would involve the recognition that the metalanguage is vague too, and therefore requires a fuzzy semantics given the fuzzy approach to vagueness. But since fuzzy semantics is supposed to invalidate classical logic (for example, the law of excluded middle), fuzzy logicians' use of classical metalogic would be illegitimate.

Fuzzy logic is a special case of many-valued logic, where logical operators act truth-functionally on more than two truth-values (degrees of truth). Many theorists have applied a three-valued logic of truth, falsity and neutrality to vagueness, although it is unclear why a three-fold classification should work if a two-fold classification does not. Higher-order vagueness is a pressing problem for many-valued logic quite generally.[4] The earlier problem of penumbral connections is also damaging for any many-valued approach with a neutral value that a sentence shares with its negation, as in standard three-valued logics. A natural response to these problems is supervaluationism.

3. Supervaluationist semantics

According to supervaluationism, a vague language admits a range of different classical assignments of referents to terms and truth-values to sentences.[5] Some but not all classical assignments are in some sense compatible with speakers' use of vague terms, the context of utterance and the extralinguistic facts; they are the *admissible valuations* (speakers themselves may be incapable of completely specifying such valuations). For example, each admissible valuation assigns an extension to 'bald', containing everybody who is definitely bald, nobody who is definitely not bald and anybody as bald as somebody whom it contains. If Jack is borderline, he is in the extension of ‘bald’ on some admissible valuations and not on others, depending on where they put the cut-off point. Each (declarative) sentence is true or false on each valuation. A borderline sentence is true on some admissible valuations, false on others.

Consider a language with the same expressions as in section 2. Supervaluationists relativize reference to a valuation V in addition to an assignment a of values to variables. If e is a variable, $\text{Ref}_{V,a}(e)$ is simply $a(e)$, the value a assigns to e, which is independent of V. If e is a constant, $\text{Ref}_{V,a}(e)$ is independent of a, but varies with V if e is vague. For example, the vague name ‘Everest’ refers to objects with slightly different spatiotemporal boundaries according to different admissible valuations. As in classical semantics, an n-place predicate refers to a function from n-tuples of objects to truth-values (truth or falsity). Since classical semantics emerges as a special case of fuzzy semantics when degrees of truth are restricted to 1 and 0 (now conceived as truth and falsity), we can write the clauses for the standard logical operators as in section 2, with the extra valuation parameter V:

SUPERatom $\quad \text{Val}_{V,a}(Ft_1...t_n) = \text{Ref}_{V,a}(F)(\text{Ref}_{V,a}(t_1),...,\text{Ref}_{V,a}(t_n))$

SUPER~ $\quad \text{Val}_{V,a}(\sim\alpha) = 1\text{-Val}_{V,a}(\alpha)$

SUPER& $\quad \text{Val}_{V,a}(\alpha \ \& \ \beta) = \min\{\text{Val}_{V,a}(\alpha), \text{Val}_{V,a}(\beta)\}$

SUPER∨ $\quad \text{Val}_{V,a}(\alpha \vee \beta) = \max\{\text{Val}_{V,a}(\alpha), \text{Val}_{V,a}(\beta)\}$

SUPER∀ $\text{Val}_{V,a}(\forall v\alpha) = \text{glb}\{\text{Val}_{V,a^*}(\alpha)$: a^* an assignment differing from a at most on $v\}$

SUPER∃ $\text{Val}_{V,a}(\exists v\alpha) = \text{lub}\{\text{Val}_{V,a^*}(\alpha)$: a^* an assignment differing from a at most on $v\}$

Here the non-classical effect of the fuzzy clauses disappears because the supervaluationist Val delivers only 0 and 1 as values. The only non-classical feature of these clauses is the parameter *V*, which so far is doing no work.

Supervaluationists use classical metalogic. Thus for each sentence α, valuation *V* and assignment *a*, either $\text{Val}_{V,a}(\alpha)$ is 1 or $\text{Val}_{V,a}(\sim\alpha)$ is 1 by SUPER~, either way, by SUPER∨, $\text{Val}_{V,a}(\alpha \vee \sim\alpha)$ is 1. Thus the law of excluded middle holds, even for borderline sentences. More generally, every theorem of classical logic is true on each valuation. Similarly, on each valuation modus ponens and other classical rules preserve truth.[6]

Supervaluationist semantics elegantly solves the problem of penumbral connections on which fuzzy semantics came to grief. If Jack and Mack are borderline for baldness, *Bj* and *Bm* are true on some admissible valuations and false on others. But since they have qualitatively identical scalps, no admissible valuation evaluates *Bj* and *Bm* differently; $\text{Val}_{V,a}(Bj)$ and $\text{Val}_{V,a}(Bm)$ are both 1 or both 0; either way, $\text{Val}_{V,a}(Bj \mathbin{\&} \sim Bm)$ is 0. Thus *Bj* & ~*Bm* is false on every valuation, as it should be; likewise for the contradiction *Bj* & ~*Bj*.

On traditional versions of supervaluationism, α is true absolutely if and only if α is true (has value 1) on every admissible valuation ('supertruth'); α is false absolutely if and only if α is false (has value 0) on every admissible valuation ('superfalsity'). Thus the principle of bivalence fails: borderline sentences are neither true nor false (absolutely), even though they satisfy the law of excluded middle. The true disjunction α ∨ ~α has no true disjunct when α is borderline. Equally, the false conjunction α & ~α then has no false conjunct. Similarly, a true existential generalization may have no true instance, for example 'For some number *n*, a man with *n* hairs is bald and a man with *n*+1 hairs is not bald'. Again, a false universal generalization may have no false instance, for example 'For every number *n*, either a man with *n* hairs is not bald or a man with *n*+1 hairs is bald'. Every valuation specifies a cut-off, but it varies across valuations.

If supervaluationists reject the principle of bivalence, they must also reject Tarskian biconditionals such as these:

TARSKI 'Jack is bald' is true if and only if Jack is bald.

'Jack is not bald' is true if and only if Jack is not bald.

For since they accept that Jack is either bald or not bald, accepting TARSKI would commit them by their classical metalogic (and their assumption that α is false if ~α is true) to the conclusion that 'Jack is bald' is either true or false, which they reject. To combine acceptance of excluded middle with rejection of bivalence they must reject TARSKI.

Supervaluationism enables us to introduce another truth predicate *T* that does satisfy Tarski's constraints (Fine 1975: 296). As usual, we give the meaning of an expression by specifying its reference relative to each valuation. Since *T* is a 1-place predicate, its referent maps objects to truth-values. Relative to any valuation *V* and assignment *a*, *T* refers to a function that maps an object o to truth if o is a true sentence of the original language relative to *V* and *a*, and otherwise maps o to falsity. Consequently, for any sentence α of the original language, the Tarskian biconditional *T*'α' ≡ α is true on any valuation. Since those biconditionals are central to the inferential role that we expect a truth predicate to play, why do traditional supervaluationists take 'supertrue' rather than *T* to express truth? The sentence

T'α' is just as vague as α itself; they vary in truth-value across admissible valuations in exactly the same way. Perhaps the original idea was to avoid using T in analysing vagueness because T is vague, and therefore a poor theoretical instrument. Supervaluationists sought to formulate a precise analysis of vagueness by using 'supertrue'. But that motivation is misconceived, for it neglects higher-order vagueness. 'Supertrue' is vague too, for 'admissible' is vague. One can readily check that by trying to see where on the colour spectrum 'That shade is red' passes from supertrue to unsupertrue, and where it passes from unsuperfalse to superfalse; it is as hopeless as trying to see where red passes into non-red. Some shades are neither definitely red nor definitely not red; sometimes 'That shade is red' is neither definitely supertrue nor definitely not supertrue. Thus the identification of truth with supertruth does not compensate for the loss of the Tarski biconditionals by any gain of precision (Williamson 1994: 162-4). More recent versions of supervaluationism therefore tend to avoid the identification, or at least to claim that 'true' is ambiguous between supertruth and Tarskian truth (McGee and McLaughlin 1995; for discussion also Andjelković and Williamson 2000).

Whether truth is supertruth or Tarskian truth, supervaluationists can characterize borderline cases within the object-language using the 'definitely' operator Δ. The idea is that Δα is true when α is true on all admissible valuations. Thus Δ functions like a universal quantifier, but varies the valuation rather than the assignment. A complication is that since 'admissible' is vague, it too must be relativized to a valuation; we write 'admissible by V'. Each valuation makes a ruling as to which valuations count as admissible. To emphasize the analogy with ∀, we state the rule for Δ thus:

SUPERΔ $\quad \mathrm{Val}_{V,a}(\Delta\alpha) = \mathrm{glb}\{\mathrm{Val}_{V^*,a}(\alpha): V^* \text{ a valuation admissible by } V\}$

$\Delta\alpha$ is true on V if α is true on every valuation admissible by V, and false on V otherwise. If we omitted the relativization 'by V' from SUPER Δ, $\Delta\alpha$ would be true on all valuations or on none, so $\Delta\Delta\alpha \vee \Delta{\sim}\Delta\alpha$ would be true on all valuations; that formula forbids borderline status to $\Delta\alpha$, thereby denying higher-order vagueness. Formally, SUPERΔ works like standard possible worlds semantics for the necessity operator $\Box$, where necessity is truth in all possible worlds. Possible worlds become admissible valuations; contingency becomes borderline status. Just as a relation of admissibility between valuations permits higher-order vagueness, so a relation of possibility ('accessibility') between worlds permits contingent necessity or possibility in modal logics weaker than the system S5. For simplicity, we ignore higher-order vagueness where it is peripheral to the discussion (but see Fine 1975: 287-98, Williamson 1994: 156-61 and 1999a).

SUPERΔ gives Δ nice features. In particular, it makes definiteness closed under logical consequence: if the argument from the premises $\alpha_1, ..., \alpha_n$ to the conclusion β is valid (truth-preserving on all valuations and assignments), so is the argument from $\Delta\alpha_1, ..., \Delta\alpha_n$ to $\Delta\beta$. Furthermore, on the reasonable assumption that every admissible valuation rules itself admissible, definiteness is factive: the argument from $\Delta\alpha$ to α is valid.

Problematic supervaluationist claims based on the identification of truth with supertruth correspond to less problematic claims about definiteness. For example, a disjunction can be definite even though no disjunct is definite; if α is borderline, $\Delta(\alpha \vee {\sim}\alpha)$ is true but $\Delta\alpha \vee \Delta{\sim}\alpha$ false. Similarly, an existential generalization can be definite even though no instance is definite; if Bn says that a man with n hairs is bald, $\Delta\exists n(Bn \;\&\; {\sim}Bn{+}1)$ is true but $\exists n\Delta(Bn \;\&\; {\sim}Bn{+}1)$ false.

Supervaluationists can even postulate a scale of degrees of truth; such degrees are not the preserve of fuzzy logic (Kamp 1975: 137-45, Williamson 1994: 154-6). For, given a suitable notion of proportion, we can define the degree of truth of α as the proportion of admissible valuations on which α is true. Thus Δα gives α degree of truth 1 and Δ~α gives it degree of truth 0. Like probabilities, and unlike degrees of truth in fuzzy semantics, supervaluationist degrees of truth falsify degree-functionality; the supervaluationist degree of truth of a conjunction or disjunction is no function of the degrees of truth of its conjuncts or disjuncts.

4. Borderline states of affairs

Vagueness is articulated by the 'definitely' operator (as in Δα) and its derivative 'it is vague whether' (as in ~ Δα & ~Δ~α). They operate on sentences (such as α). Vagueness concerns borderline cases; in a borderline case (for example, when it is vague whether Jack is bald) there is a judgement to be hesitated over (the judgement that Jack is bald); the expression of a judgement is in a sentence ('Jack is bald'). For simplicity, let us suppose that expressions of different grammatical categories are all correlated with different elements of reality, their *ontological correlates*. The ontological correlate of a sentence is a state of affairs, just as the ontological correlate of a singular term is an object and the ontological correlate of a predicate is a property or relation. We should therefore expect the question of vagueness in reality to arise primarily for states of affairs, not for objects or even properties and relations. 'Is reality vague?' is not to be paraphrased as 'Are there vague objects?' or 'Are there vague properties and relations?'.

What is a state of affairs? For any object o and any property P, there is the state of affairs that o has P; it obtains if and only if o has P. For any objects o_1 and o_2 and any binary relation R, there is the state of affairs that o_1 has R to o_2; it obtains if and only if o_1 has R to o_2.[7] There are doubtless many other kinds of states of affairs. For present purposes, we need not take the ontology of states of affairs wholly seriously. Once we have the items that determine them (such as objects, properties and relations), we could speak of those items directly. But the notion of a state of affairs is convenient, because it enables us to generalize without surveying the items that compose states of affairs.

States of affairs can be individuated in coarse-grained or fine-grained ways. On a fine-grained conception, states of affairs are structured items, with objects, properties and relations as constituents; S and S* may be distinct states of affairs because they are differently constituted even if, necessarily, S obtains if and only if S* obtains. On a less fine-grained conception, S and S* are identical if and only if, necessarily, S obtains if and only if S* obtains. Such states of affairs might be identified with classes of possible worlds. Formally, it would be simplest to use a *very* coarse-grained conception, on which states of affairs are simply truth-values or degrees of truth, for those are the items correlated with sentences by the extensional semantic theories in sections 2 and 3. Truth obtains and falsity does not; a degree of truth obtains to that very degree. Although the phrase ‘state of affairs’ becomes rather misleading when applied to truth-values or degrees of truth, what matters is that sentences stand to them in a relation something like reference. Once modal and temporal operators are introduced, sentences might be treated as referring to functions from possible worlds and times to truth-values or degrees of truth; such functions are more naturally conceived as formal versions of states of affairs. In any case, the arguments below do not require the coarse-grained conception; they run on any reasonable standard of individuation.

Since states of affairs are the ontological correlates of sentences, the most natural way to generalize over states of affairs is by quantifying into sentence position. Indeed, in a metaphysically deeper discussion, we might *explain* apparent quantification over states of affairs as a crude rendering into a natural language of quantification into sentence position. For example, to say that some state of affairs does not obtain, we write $\exists S{\sim}S$. In English, we have the noun phrase ‘state of affairs’ and must use the verb ‘obtain’ to construct a corresponding sentence; but that is an artifact of the difficulty of expressing quantification into anything but noun phrase position unambiguously in natural language. The formal language achieves the desired effect more economically. This quantification into sentence position should not be understood substitutionally. Just as some objects may lack names, so some states of affairs may lack sentences to express them. Rather, the quantification should be understood in the normal way, by variation in the assignment to the relevant variable (‘*S*’). Any state of affairs may be assigned to a variable in sentence position.

‘...’ expresses a borderline case if and only if it is vague whether .... We therefore define a state of affairs S to be borderline if and only if it is vague whether S obtains. Reality is vague if and only if at least one state of affairs is borderline. For if reality is vague, it is vague how things are, so for some way it is vague whether things are that way; thus, for some state of affairs S, it is vague whether S obtains. Conversely, if for some state of affairs S it is vague whether S obtains, for some way it is vague whether things are that way, so it is vague how things are; thus reality is vague. Reality is precise if and only if it is not vague: no state of affairs is borderline, every one either definitely obtains or definitely fails to obtain.[8] Again, the formal language expresses the idea more economically. To say that reality is vague, that some state of affairs neither definitely obtains nor definitely fails to obtain, we write:

(1) $\exists S({\sim}\Delta S \ \&\ {\sim}\Delta{\sim}S)$

To say that reality is precise, we write the negation of (1), or equivalently $\forall S(\Delta S \vee \Delta{\sim}S)$.

5. Supervaluationist states of affairs

Does the account in section 4 trivialize the claim that reality is vague? Everyone agrees that it is vague whether this [I point] is a heap. Thus it is vague whether the state of affairs that this is a heap obtains. Therefore, apparently, of at least one state of affairs (that this is a heap) it is vague whether it obtains, so at least one state of affairs is borderline, so reality is vague. That conclusion looks cheap; the case is paradigmatically one in which we want to locate the vagueness in our words or concepts.

The argument is fallacious. It moves from 'It is vague whether the state of affairs that this is a heap obtains' to 'Of the state of affairs that this is a heap, it is vague whether it obtains'. That is no more valid than the fallacious move from 'It is contingent whether the number of planets is even' to 'Of the number of planets, it is contingent whether it is even' (it is contingent how many planets there are, but not whether a given number is even). In the premise, a definite description ('the state of affairs that this is a heap', 'the number of planets') is in the scope of an operator ('it is vague whether', 'it is contingent whether'); in the conclusion their scopes have been illegitimately reversed. Just as for no number n is it necessary that there are exactly n planets, so perhaps for no state of affairs S is it definite that S obtains if and only if this is a heap. Thus reality may be precise even though it is vague whether this is a heap.

Supervaluationism realizes this possibility in a natural way. It is vague whether this is a heap (Ht):

(2) $\sim\Delta Ht \;\&\; \sim\Delta{\sim}Ht$

The question is whether (2) entails (1), perhaps by existential introduction. As quantification into sentence position is comparatively unfamiliar, we may start by assessing separate existentially generalizations into name and predicate position in (2):

(1*) $\exists X\exists x(\sim\Delta Xx \;\&\; \sim\Delta{\sim}Xx)$

According to (1*), for some property and object it is vague whether the latter has the former. It will emerge that, under supervaluationism, (1*) is false and does not follow from (2). By an extension of the argument, (1) is false and does not follow from (2).

First consider (2). Since *H* and *t* are constants rather than variables, we can ignore the assignment *a*. In this context we may treat the singular term *t* ('this') as precise; for some object o, its referent $\mathrm{Ref}_{V,a}(t)$ is o for every admissible valuation *V*. But the predicate *H* ('is a heap') is vague; its extension contains o on some but not all admissible valuations. For some admissible valuations *V*, by SUPERatom, $\mathrm{Val}_{V,a}(Ht) = \mathrm{Ref}_{V,a}(H)(\mathrm{o}) = 1$, so $\Delta{\sim}Ht$ is false. For other admissible valuations *V*, $\mathrm{Val}_{V,a}(Ht) = \mathrm{Ref}_{V,a}(H)(\mathrm{o}) = 0$, so ΔHt is false. Thus (2) is true. That is the standard supervaluationist treatment of a borderline case: the sentence varies in truth-value across admissible valuations.

Now take (1*). Consider the truth-value of *Xx*. By SUPERatom, $\mathrm{Val}_{V,a}(Xx) = \mathrm{Ref}_{V,a}(X)(\mathrm{Ref}_{V,a}(x))$. Since *X* and *x* are variables, their referents are just what *a* assigns to them; $\mathrm{Ref}_{V,a}(X) = a(X)$ and $\mathrm{Ref}_{V,a}(x) = a(x)$. Consequently, $\mathrm{Val}_{V,a}(Xx) = a(X)(a(x))$. *Xx* is true if the object assigned to *x* is in the extension assigned to *X* and false otherwise. The crucial point is that $a(X)(a(x))$ depends on *a* but not on *V*. Consequently, for any given assignment,

Xx is true on all valuations or none. Thus $\sim\Delta Xx$ & $\sim\Delta\sim Xx$ is false on all assignments, so by SUPER∃ (1*) is false too. Thus (1*) does not follow from (2). On this treatment, it cannot be vague whether an object has a property.

The argument against (1*) works because variables are not vague. It therefore extends to an argument against (1). The simplest extension treats *S* as a 0-place predicate variable. By SUPERatom, $\mathrm{Val}_{V,a}(S)$ would be $a(S)(\langle\rangle)$, where $a(S)$ is a function from 0-tuples (of which there is only one) to truth-values and <> is the 0-tuple; this is tantamount to having *a* assign *S* a truth-value. There are more complex possibilities too. What matters is that there is no more room for the truth-value of *S* to depend on the valuation *V* than there was for the truth-value of *Xx* to do so, because the assignment alone fixes the value of the variable. For any given assignment, *S* is true on all valuations or none. Thus $\sim\Delta S$ & $\sim\Delta\sim S$ is false on any assignment, so (1) is false and $\forall S(\Delta S \vee \Delta\sim S)$ true. Thus (1) does not follow from (2). On this supervaluationist treatment, it cannot be vague whether a state of affairs obtains. Reality is guaranteed to be precise.

Under supervaluationism, one can existentially generalize into the scope of Δ only if the expression on which one generalizes refers precisely. For the singular term *t*, the condition is $\exists x\Delta x{=}t$; for the predicate *H*, it is $\exists X\Delta\forall x(Xx \equiv Hx)$; for the sentence *Ht*, it is $\exists S\Delta(S \equiv Ht)$. Those conditions for *H* and *Ht* are not met here. Supervaluationists grant the weaker conditions $\Delta\exists X\forall x(Xx \equiv Hx)$ and $\Delta\exists S(S \equiv Ht)$, but they do not suffice for existential generalization.

The key to the falsification of (1) and (1*) is that variables are precise: on a given assignment, their reference is constant across valuations. One might therefore suppose that supervaluationists could verify (1) and (1*) simply by having assignments assign values to variables relative to valuations. But that change collapses a distinction crucial to supervaluationists. They admit that, definitely, there is a cut-off point for a vague predicate *F* ($\Delta\exists n(Fn$ & $\sim Fn+1)$), since every admissible valuation has such a point. But they insist that *F* is still vague because no point is definitely the cut-off ($\sim\exists n\Delta(Fn$ & $\sim Fn+1)$); the point varies across valuations. If the assignment of values to variables were relativized to valuations, $\Delta(Fn$ & $\sim Fn+1)$ would be true when, relative to each valuation, the variable *n* was assigned the cut-off number for that valuation; thus, by SUPER∃, $\exists n\Delta(Fn$ & $\sim Fn+1)$ would be evaluated as true and the supervaluationist claim of indefiniteness would disappear. The precision of variables is no technical accident. It is crucial to supervaluationists' articulation of their main idea. To focus on states of affairs rather than our ways of referring to them, we generalize with variables whose reference is fixed to one states of affairs across valuations; in doing so, we eliminate vagueness. If vagueness is variation across admissible valuations then reality is precise.

Many theorists of vagueness accept the metaphysics of supervaluationism but reject its semantics. Like supervaluationists, they conceive vague thought and language as related only indirectly to an underlying reality (typically, one described by popular physics), but disagree on the semantic consequences of that shared conception. Nihilists, for instance, take vague terms to suffer reference failure, and therefore vague predications to be truth-valueless or false; although they may share the supervaluationists' view of what there is to be referred to, they take a harsher view of the semantic consequences of indecision between potential referents. Under nihilism too, reality is precise.[9]

6. Fuzzy states of affairs

Section 4 explicated the claim that reality is vague. Section 5 asked whether that explication makes the claim trivially true, and answered ‘No’ by showing that it makes the claim false given a standard form of supervaluationism. But that easy argument raises the opposite question: does the explication make the claim trivially false? By showing that it makes the claim true given a standard form of fuzzy logic, this section will answer ‘No’ to that question too. On whether reality is vague, supervaluationism and fuzzy logic stand opposed.

As before, the issue is whether a banal statement of a borderline case, for example (2), that it is vague whether this is a heap, entails (1), that reality is vague. And, as before, it is easiest to start by asking whether (2) entails the more specific claim (1*), that for some property and some object, it is vague whether the latter has the former. Can we existentially generalize into the scope of ‘it is vague whether’ and therefore of Δ? The semantics of those operators is more straightforward under fuzzy logic than under supervaluationism. Indeed, Δ is as much of a degree-function (the analogue of a truth-function) as negation is on the fuzzy semantics; quantifying into the scope of Δ is no more problematic than quantifying into the scope of negation. In classical semantics, the predicate variable X can be assigned any function from objects to truth-values; likewise, in fuzzy semantics, it can be assigned any function from objects to degrees of truth, in particular that to which the predicate H refers. Similarly, the variable x can be assigned any object, in particular that to which the singular term t refers. Thus by FUZZY$\exists$ (1*) is a straightforward consequence of (2). For analogous reasons, (1) is also a straightforward consequence of (2). The extension of the argument is simplest if we regard S as a 0-place predicate variable, but holds on more complex interpretations too. In fuzzy logic, *every* borderline case, however ‘linguistic’ in appearance, makes reality vague.

At first sight the result is disconcerting. On reflection, however, it presents a problem for fuzzy logic, not for the explication of the claim that reality is vague. Given FUZZYΔ, the degree of truth of $\Delta\alpha$ depends only on the degree of truth of α, just as the degree of truth of $\sim\alpha$ depends only on the degree of truth of α; $\Delta\alpha$ is just as much about what α is about as $\sim\alpha$ is. Δ does not represent any kind of semantic ascent to a metalinguistic level. Thus vagueness in whether this is a heap is straightforward vagueness in how things are. By contrast, SUPERΔ makes Δ a kind of universal quantifier over admissible valuations, which is a kind of semantic ascent. Of course, it is natural to object to FUZZYΔ that it makes vagueness in reality come far too cheap. Perhaps there is no single state of affairs that this is a heap but many states of affairs concerning the exact number and arrangement of grains. But that suggestion is more congenial to supervaluationism than to fuzzy logic. For FUZZYatom relates 'This is a heap' to reality as directly as it does sentences concerning the exact number and arrangement of grains. To take the fuzzy semantics at face value is to treat vagueness in whether this is a heap as simply vagueness in how things are. If it is not vagueness in how things are, then something is wrong with the fuzzy semantics.

According to supervaluationism, no borderline case makes reality vague. According to fuzzy semantics, any borderline case makes reality vague. Of course, there will be no end of mixed views, perhaps embodying elements of both supervaluationism and fuzzy semantics, according to which some but not all borderline cases make reality vague. Some may even insist on applying 'supervaluationism' or 'fuzzy logic' as a label to such a mixed view. We cannot survey all the possible combinations here. Nevertheless, we may wonder whether an account of vagueness can distinguish in any principled way between some borderline cases that make reality vague and others that do not. At any rate, it has become obvious that the question 'Is reality vague?' must eventually be answered by comparing theories of vagueness overall.

7. Vague objects, properties and relations[10]

'Is reality vague?' and 'Are there vague objects?' are often treated as the same question. That conflation is symptomatic of a tendency to conceive reality as merely a collection of objects, as though one could describe it fully by listing them, without having to specify their properties and relations. The tendency has been elevated to an explicit claim, the truthmaker principle, according to which (in its unqualified form) for every truth an object exists whose existence is sufficient for that truth. Once we take quantification into predicate and sentence position seriously in its own right, the truthmaker principle looks unmotivated, for it assigns a metaphysically privileged status to quantification into name position (over objects). It derives ontology from linguistic prejudice (Williamson 1999b).

Although 'Are there vague objects?' is not equivalent to 'Is reality vague?', the former question might be interesting on its own terms. What could it mean to call an object 'vague'? Objects are the referents of names and other singular terms; vagueness is characterized by operators on sentences. Of course, a word is an object and can be vague in the ordinary sense; but it is vague because there are or could be cases in which it is vague whether it applies. 'It applies' is a sentence, not a singular term, and application has no obvious analogue for objects other than expressions and concepts.

In many ordinary cases it is tempting to say that there is a vague object. Where, for example, are the boundaries of Mount Everest? Of some spatial points, it is vague whether Everest includes them; for some rocks, it is vague whether they are part of Everest. We might conclude that Everest is a vague object. But what, if anything, is special about the relations of location and parthood? And what justifies the move from vagueness expressed at the level of the sentences ‘Everest includes those points’ and ‘Those rocks are part of Everest’ to vagueness expressed at the level of the name ‘Everest’?

Supervaluationists can significantly ascribe vagueness to referring expressions of any syntactic category. The test is whether they vary in reference across admissible valuations. Thus the name ‘Everest’ is vague if on two such valuations it refers to different mountainous objects. But that is not vagueness in the objects referred to. As argued in section 5, supervaluationism implies that it is never vague how things are. For example, either ‘is part of’ is vague or it is not. If it is vague, for some objects o and o* it is vague whether o is part of o*, but we are not justified in attributing the vagueness to o or o*. If ‘is part of’ is precise, then for no objects o and o* is it vague whether o is part of o*. Either way, under supervaluationism, we have no basis for classifying objects as vague.

What of fuzzy semantics? ‘Is part of’ refers to a function that maps the ordered pair of Everest and a peripheral rock to an intermediate degree of truth. On the face of it, that is vagueness in the function, not in the objects. The fuzzy approach seems to provide a natural notion of vagueness for the referents of predicates (which we may call ‘properties’ and ‘relations’), as the capacity to yield intermediate degrees for some objects, but no natural notion of vagueness for objects, the referents of singular terms. Could fuzzy logicians define an object o to be vague if and only if for some object o* it is vague whether o* is part of o (o* is part of o to an intermediate degree)? That would be a more or less arbitrary stipulation, without natural grounding in the fuzzy semantics. Why should vagueness in whether o* is part of o be attributed to o rather than to o* or to parthood? Indeed, the admission of vague objects would undermine the original conception of vague properties and relations, for why should vague objects not have precise properties (such as having mass exactly *m*) or relations (such as having exactly the same mass) to intermediate degrees? Even granted a notion of a precise object, we cannot determine whether a property is vague by asking whether it applies to precise objects to intermediate degrees, for the degree to which any precise object has the intuitively vague property of being a very vague red object is 0. Such difficulties for the conception of vague objects generalize beyond fuzzy logic.

We have no good reason to enter the maze created by the conception of vague objects. It depends on the attempt to attribute vagueness at the level of subsentential expressions. Without a theory of vagueness, we should not assume that the attempt is sensible. After all, it would be foolish to attribute *falsity* at the level of subsentential expressions, to ask whether the falsity of 'Cats bark' should be blamed on falsity in 'cats' or cats or on falsity in 'barks' or barking. We need some reason to suppose that vagueness distributes in a way in which falsity does not. Fuzzy logic supplies no such reason. Although supervaluationism provides a reason, its way of distributing vagueness amongst subsentential expressions forbids its projection onto their referents. Such considerations suggest that we should abandon the question 'Are there vague objects?'. We can do so the more easily because we do not thereby abandon the question 'Is reality vague?'.

8. Vague identity

Much of the literature on vague objects focusses on whether it can be vague of objects whether they are identical. Problem cases of identity might suggest a positive answer. If it is unclear whether the person who emerges from drastic brain surgery is the same as the one who was pushed into the operating theatre, perhaps the earlier person and the later one are objects of which it is vague whether they are identical. However, Gareth Evans (1978) proposed a general argument to show that it cannot be vague whether objects are identical. In brief, suppose that it is vague whether o is o*; since it is not vague whether o* is o*, something holds of o that does not hold of o* (that it is vague whether it is o*); but Leibniz's Law of identity states that if o is o* then whatever holds of o holds of o*; therefore o is not o*.[11] That is not yet a straight contradiction. However, by an extension of the argument, from the definiteness of its premise we can infer the definiteness of its conclusion: given that it is *definitely* vague whether o is o*, it follows that *definitely* o is not o*, and therefore that it is not vague whether o is o*. But given that it is definitely vague whether o is o*, it also follows that it *is* vague whether o is not o*. Thus at least the supposition that it is definitely vague whether o is o* seems to yield a contradiction. Many defenders of vagueness in reality or of vague objects have felt obliged to find a fallacy in Evans's argument, although it is not obvious why the more general issues should be thought to turn on the very special relation of identity.[12]

For most supervaluationists, '=' has the same extension on every admissible valuation; it is uniquely identified by its structural characteristics (its extension contains the ordered pair of each object with itself and is a subclass of every class with that property). Thus the sentence $t=t^*$ varies in truth-value across admissible valuations only if at least one of the singular terms t and t^* varies in reference across such valuations. Consequently, we cannot quantify into $\sim\Delta t=t^*$ & $\sim\Delta\sim t=t^*$ to conclude that for some objects it is vague whether they are identical (Thomason 1982). That is consistent with Evans's idea that the formula defeats itself on the assumption that t and t^* refer precisely to objects between which identity is vague (Lewis 1988). His intended conclusion is supervaluationistically correct, as the argument of section 5 implies.

Evans's argument is more interesting in a fuzzy context, where existential generalization holds. Suppose that, relative to some assignment, $x=x^*$ is true to some intermediate degree. Thus $\sim\Delta x=x^*$ is true to degree 1. But $x=x$ is also true to degree 1, as therefore is $\Delta x=x$. Consequently, the argument from $\Delta x=x$ and $\sim\Delta x=x^*$ to $\sim x=x^*$ has premises true to degree 1 and a conclusion true to a degree less than 1: it does not preserve definite truth in fuzzy logic, even though it has the form of an argument from $\varphi(x)$ and $\sim\varphi(x^*)$ to $\sim x=x^*$. Strictly speaking, what fails is not Leibniz's Law (from $x=x^*$ and $\varphi(x)$ to $\varphi(x^*)$) but a contraposed variant (Parsons 1987). Nevertheless, the contraposed variant is fundamental to our thinking about identity. If what we know is that this man is of blood group O and that John is not of blood group O, how else are we to conclude that this man is not John? Moreover, even the uncontraposed Leibniz's Law fails to preserve intermediate degrees of truth. Under the same supposition as before, the argument from $x=x^*$ and $\Delta x=x$ to $\Delta x=x^*$ has both premises true to a degree greater than 0 and a conclusion true to degree 0.

The natural suspicion is that fuzzy logicians permit degrees of identity only by losing their grip on the notion of identity. Consider the classical metalogic on which fuzzy logicians rely. Since they take vague terms to require nonclassical logic, they are treating the metalanguage as precise, its sentences true to degree 1 or 0. Since '=' figures in the metalanguage, $x=x^*$ is true in the metalanguage to degree 1 or 0 relative to any assignment. Unless $x=x^*$ is true in the object-language to the same degree relative to the same assignment, '=' is not being interpreted as identity. Consequently, fuzzy logicians should assign intermediate degrees of truth to identity sentences only if they adopt a non-classical metalogic. The same goes for other forms of many-valued logic.

Opponents of Evans's argument often claim to show that it is fallacious or question-begging by constructing many-valued models that invalidate one or more of its steps within the framework of a classical metalogic (van Inwagen 1988, Parsons and Woodruff 1995, Priest 1998). The foregoing considerations show this method to be unsound, for it depends on interpreting '=' in the object-language as meaning something other than identity. After all, we can construct a formal model in which $t=t^*$ and Ft are true and Ft^* is false by interpreting '=' to mean distinctness. To claim on those grounds that Leibniz's Law is fallacious or question-begging would be silly. The purported many-valued countermodels to Evans's argument make the same mistake, albeit in a far subtler form (Williamson 2002). No adequate treatment of Evans's argument within a systematically fuzzy or many-valued metalogic has been provided. Moreover, the earlier problem of penumbral connections gives us ample reason to reject approaches based on fuzzy or many-valued logic.

9. Epistemicism

Section 3 argued that supervaluationists' rejection of Tarskian constraints on truth and falsity is poorly motivated. But if they accept those constraints, they must also accept the principle of bivalence: 'This is a heap' is either true or false, although we have no way of knowing which. Thus supervaluationism threatens to collapse into an epistemic view of vagueness (Williamson 1995, 1997a: 216-17). Rather than examine the putative collapse, let us consider epistemicism in its own right. On this view, truth is Tarskian and logic is classical. Vagueness consists in a special kind of irremediable ignorance in borderline cases.[13]

What prevents us from knowing whether this is a heap? Suppose that it is in fact a heap. I might judge, truly, 'This is a heap'. But, in a borderline case, I could easily have judged 'This is a heap' even if the factors that determine the reference of my word 'heap' had differed slightly, making me express a different and false proposition. In that sense I cannot discriminate the counterfactual assignment of reference to my words from the actual assignment; the two semantic valuations are indiscriminable for me. Our powers of discrimination are limited; we cannot make our judgements perfectly sensitive to all the reference-determining factors. Thus my actual judgement is unreliably based; even if true, it does not constitute knowledge. That, in outline, is an epistemicist explanation of our irremediable ignorance.

At least in its simplest form, epistemicism can take over and reinterpret the formal apparatus of supervaluationism, with a Tarskian conception of truth. To say that a valuation V^* is admissible by a valuation V is now to say that V^* is indiscriminable from V in the sense indicated. Since Δ now has an epistemic sense, we can read it as 'clearly' rather than 'definitely'.

If epistemicism has the same formal structure as supervaluationism, it too implies that reality is precise. Suppose that it is vague whether this rock is part of Everest and that 'this rock' and 'is part of' are precise. Then both supervaluationists and epistemicists deny that Everest is an object of which it is vague whether this rock is part of it. They assert that Everest is an object of which it is vague whether it is Everest (for no object is the referent of 'Everest' on all admissible valuations), although it is not vague whether Everest is Everest. For epistemicists, this simply reflects our limited capacity to discriminate between situations in which the name 'Everest' refers to Everest and situations in which it refers to entities that coincide with Everest only approximately; in particular, some differ from Everest in whether they include this rock. Yet when we ask ourselves 'Is this rock part of Everest?', 'Everest' refers to a unique object. Consider an analogy in which we prescind from vagueness. Let the name 'Es' refer to England if there is life in other galaxies and to Scotland otherwise. We know that Cumbria is part of England and not part of Scotland. Not knowing whether there is life in other galaxies, we know neither that Es is England nor that Es is Scotland. We do not know whether Cumbria is part of Es. Yet when we ask ourselves 'Is Cumbria part of Es?', 'Es' refers to a unique object, perhaps England, perhaps Scotland.

The epistemic point that no object is clearly Everest does not at all undermine the metaphysical point that some object is uniquely Everest. By contrast, although supervaluationists grudgingly admit that, definitely, some object is uniquely Everest, they take their admission to be somehow metaphysically unserious because no object is definitely Everest; no object is definitely the unique referent of 'Everest'. Supervaluationists conceive their relativization of reference to valuations as making the relation between language and reality somehow less direct than it is in classical truth-conditional semantics. For epistemicists, the relativization has no such effect: it is simply a technical device within classical truth-conditional semantics to handle an epistemic operator. Consequently, epistemicism does not involve the second-guessing of our vague metaphysical beliefs that is characteristic of supervaluationism.

How can it be epistemically vague which object is Everest? Surely a mountaineer can look at Everest, say 'That mountain is Everest', and thereby express knowledge, of Everest, that it is Everest. Visual acquaintance provides a paradigm of such *de re* knowledge. But if it is known, of Everest, that it is Everest, how can it also be epistemically vague, of Everest, that it is Everest? In situations indiscriminable from the actual one, the visual demonstrative 'that mountain' and the name 'Everest' refer to something slightly different from Everest, but still maintain coreference; in saying 'That mountain is Everest' the mountaineer would still express a truth, although a slightly different one. The *de dicto* knowledge that the mountain is Everest is unthreatened. Moreover, *de dicto* knowledge entails *de dicto* clarity; it is not epistemically vague whether the mountain is Everest. But *de re* knowledge may not entail *de* *re* clarity. Unlike *de re* clarity, *de re* knowledge may not require perfect discrimination of the object. The notion of knowing something of an object has everyday uses because we are willing to apply it to subjects who do not meet a perfectionist standard for discriminating the object. By contrast, the Δ operator is a theoretical instrument for analysing vagueness; in that project, slight failures of discrimination are just what we are interested in. Thus it can be epistemically vague of Everest whether it is Everest even though it is known of Everest that it is Everest. Our self-attribution of *de re* knowledge in such borderline cases helps to explain our temptation to attribute the vagueness to the objects themselves.

Under epistemicism, each state of affairs either clearly obtains or clearly fails to obtain. One might therefore conclude that the source of any vagueness is our way of conceptualizing the state of affairs, not the state of affairs itself. For example, suppose that it is vague whether 17 is small (for a natural number). For some natural number *n*, 'small' refers in this context to the property of being less than *n*. Some numeral, say '18', refers to *n*. On some reasonable conceptions of states of affairs, the state of affairs that 17 is small is the state of affairs that 17 is less than 18. It is not vague whether 17 is less than 18. That it is vague whether 17 is small shows that the state of affairs that 17 is small *can* be conceptualized as borderline, not that it *must* be.

Are all borderline cases like that, given epistemicism? That every state of affairs can be conceptualized precisely is not obvious. Perhaps heaphood cannot be conceptualized precisely. But complete precision may be unnecessary. Suppose that 'This is a heap' (Ht) is borderline. Define artificial predicates $H+$ and $H-$ to have the same extension as H in all non-actual worlds; in the actual world, $H+x$ is equivalent to $Hx \vee x=t$ and $H-x$ to $Hx \,\&\, {\sim}x=t$. If Ht is true, $H+$ is necessarily coextensive with H, so Ht and $H+t$ arguably express the same state of affairs; but $H+t$ ($Ht \vee t=t$) is clearly true. If Ht is false, $H-t$ is necessarily coextensive with

H, so *Ht* and *H-t* arguably express the same state of affairs; but *H-t* (*Ht* & ~*t*=*t*) is clearly false. Since *Ht* is either true or false, the state of affairs that this is a heap can be conceptualized as non-borderline, unless states of affairs are so fine-grained that the additional logical structure in the new predicates makes them express new states of affairs.

Perhaps any state of affairs that can be conceptualized as borderline can also be conceptualized as non-borderline. But that might require highly artificial methods of conceptualization. Some states of affairs may be 'naturally' conceptualizable only as borderline. That property would distinguish those states of affairs from others, and it might reflect underlying intrinsic differences between the states of affairs.

Identity states of affairs illustrate the point (see also Williamson 2002). Suppose that it is vague whether Everest is Schmeverest. If the identity holds, 'Everest is Schmeverest' expresses the same state of affairs as 'Everest is Everest'; it is clear that Everest is Everest. But if the identity fails, the two sentences express different states of affairs; perhaps the state of affairs that Everest is Schmeverest cannot be 'naturally' conceptualized as non-borderline.

Epistemicism provides only an etiolated sort of metaphysical vagueness, constitutively dependent on thinkers' epistemological limitations. That may be as much metaphysical vagueness as we have any reason to expect. Common sense insists that it is vague whether this rock is part of Everest, that this rock and Everest are genuine objects and that parthood is a genuine relation; it leaves the underlying nature of the vagueness for theory to determine. Under epistemicism, common sense may be straightforwardly correct.[14]

Notes

1 For vagueness in general see Graff and Williamson 2002, Keefe and Smith 1996, Keefe 2000 and Williamson 1994.

2 Goguen 1969 is an early example of this approach. See also Forbes 1983 and 1985: 164-74.

3 For simplicity, we ignore other relativizations: to circumstances of evaluation (possible worlds and times) for modal and temporal operators, to contexts of utterance for indexicals. These complications would not harm the argument of this chapter. Possibility and time are dimensions orthogonal to vagueness. Context-shifting has been argued to play a major role in one of the main manifestations of vagueness, sorites paradoxes (Kamp 1981), but pragmatic effects are not of primary interest for the metaphysics of vagueness.

4 Tye 1990, 1994 and Horgan 1994 discuss higher-order vagueness in a three-valued context, but their remarks neither constitute a systematic metalogical treatment nor show how to construct one. See also Williamson 1994: 127-31.

5 Fine 1975 and Kamp 1975 exemplify this approach; for discussion see also Williamson 1994: 142-64 and Keefe 2000.

6 Some versions of supervaluationism do not validate classical rules involving the

discharge of premises, such as conditional proof and reductio ad absurdum (Fine 1975: 290, Williamson 1994: 150-3).

7 An ordering of the argument places of R is assumed.

8 For the contrasting proposal that the world is vague if and only if vague matters do not supervene on precise ones see Peacocke 1981: 132-3, criticized by Hyde 1998; for relevant discussion see Williamson 1994: 201-4 and McLaughlin 1997. Sainsbury 1995 finds no substance to the question whether the world is vague; he looks in the wrong place.

9 For nihilist arguments of varying degrees of extremism see Heller 1988, 1990, Horgan 1995, Unger 1979a, 1979b, 1979c, 1980 (contrast 1990: 321-3), Wheeler 1975, 1979. For discussion see Abbott 1983, Grim 1982, 1983, 1984, Sanford 1979, Williamson 1994: 165-84.

10 For some discussion of relevant issues see Akiba 2000, Burgess 1990, Lewis 1993, Lowe 1995, Quine 1981, Rolf 1980, Sainsbury 1989, 1995, Sanford 1993, Sorensen 1998, Tye 1990, 1995, 1996.

11 For a similar argument see Salmon 1982: 243-6, 1984, 1986: 110-14.

12 On vague identity see also Broome 1984, Burgess 1989, 1990, Cook 1986, Copeland 1995, 1997, Cowles 1994, Cowles and White 1991, Engel 1991: 196-8, 213-15,

French and Krause 1995, Garrett 1988, 1991, Gibbins 1982, Hawley 1998, Heck 1998, Hirsch 1999, Howard-Snyder 1991, Johnsen 1989, Keefe 1995, Lowe 1994, 1997, McGee and McLaughlin 1995, Noonan 1982, 1984, 1990, 1995, Over 1984, 1989, Parsons 1987, 2000, Parsons and Woodruff 1995, Pelletier 1984, 1989, Priest 1991, 1998, Rasmussen 1986, Stalnaker 1988, Tye 1990, van Inwagen 1988, 1990, Wiggins 1986, Williamson 1996a, 2002, Zemach 1991.

13 For development and defence of epistemicism, and references to earlier epistemicist and anti-epistemicist writings see Williamson 1994, 1995, 1996b, 1997a, 2002. Keefe 2000 contains a recent critique of epistemicism. For technical issues relevant to both epistemicism and supervaluationism see Andjelković and Williamson 2001 and Williamson 1999a.

14 Thanks to Dorothy Edgington, Delia Graff, Dean Zimmerman and an audience in Oxford for comments on an earlier draft.

Bibliography

Abbott, W. 1983. 'A note on Grim's sorites argument'. *Analysis* 43: 161-4.

Akiba, K. 2000. 'Vagueness as a modality'. *Philosophical Quarterly* 50: 359-70.

Andjelković, M., and Williamson, T. 2001. 'Truth, falsity and borderline cases'. *Philosophical Topics* 28: 211-244.

Broome, J. 1984. 'Indefiniteness in identity'. *Analysis* 44: 6-12.

Burgess, J. A. 1989. 'Vague identity: Evans misrepresented'. *Analysis* 49: 112-19.

Burgess, J. A. 1990. 'Vague objects and indefinite identity'. *Philosophical Studies* 59: 263-87.

Cook, M. 1986. 'Indeterminacy of identity'. *Analysis* 46: 179-86.

Copeland, B. J. 1995. 'On vague objects, fuzzy logic, and fractal boundaries'. *Southern Journal of Philosophy*, sup. vol. 33t: 83-96.

Copeland, B. J. 1997. 'Vague identity and fuzzy logic'. *Journal of Philosophy* 94: 514-34.

Cowles, D. 1994. 'On van Inwagen's defense of vague identity'. *Philosophical Perspectives* 8: 137-58.

Cowles, D., and White, M. 1991. 'Vague objects for those who want them'. *Philosophical Studies* 63: 203-16.

Engel, P. 1991. *The Norm of Truth*. London: Harvester Wheatsheaf.

Evans, G. 1978. 'Can there be vague objects?'. *Analysis* 38: 208.

Fine, K. 1975. 'Vagueness, truth and logic'. *Synthese* 30: 265-300.

Forbes, G. 1983. 'Thisness and vagueness'. *Synthese* 54: 235-59.

Forbes, G. 1985. *The Metaphysics of Modality*. Oxford: Clarendon Press.

French, S., and Krause, D. 1995. 'Vague identity and quantum indeterminacy'. *Analysis* 55:

20-6.

Garrett, B. 1988. 'Vagueness and identity'. *Analysis* 48: 130-4.

Garrett, B. 1991. 'Vague identity and vague objects'. *Noûs* 25: 341-51.

Gibbins, P. 1982. 'The strange modal logic of indeterminacy'. *Logique et Analyse* 25: 443-6.

Goguen, J. 1969. 'The logic of inexact concepts'. *Synthese* 19: 325-73.

Graff, D., and Williamson, T. (eds.) 2002. *Vagueness*. Aldershot: Ashgate.

Grim, P. 1982. 'What won't escape sorites arguments'. *Analysis* 42: 38-43.

Grim, P. 1983. 'Is this a swizzle stick which I see before me?'. *Analysis* 43: 164-6.

Grim, P. 1984. 'Taking sorites arguments seriously: some hidden costs'. *Philosophia* 14: 251-72.

Hawley, K. 1998. 'Indeterminism and indeterminacy'. *Analysis* 58: 101-6.

Heck, R. 1998. 'That there might be vague objects (so far as concerns logic)'. *Monist* 81: 274-96.

Heller, M. 1988. 'Vagueness and the standard ontology'. *Noûs* 22: 109-31.

Heller, M. 1990. *The Ontology of Physical Objects: Four-Dimensional Hunks of Matter*. Cambridge: Cambridge University Press.

Hirsch, E. 1999. 'The vagueness of identity'. *Philosophical Topics* 26: 139-58.

Horgan, T. 1994. 'Robust vagueness and the forced-march sorites paradox'. *Philosophical Perspectives* 8: 159-88.

Horgan, T. 1995. 'Transvaluationism: a Dionysian approach to vagueness'. *Southern Journal of Philosophy* 33: 97-126.

Howard-Snyder, F. 1991. '*De re* modality entails *de re* vagueness'. *Pacific Philosophical Quarterly* 72: 101-12.

Hyde, D. 1998. 'Vagueness, ontology and supervenience'. *Monist* 81: 297-312.

Johnsen, B. 1989. ‘Is vague identity incoherent?’. *Analysis* 49: 103-12.

Kamp, H. 1975. ‘Two theories about adjectives’. In E. Keenan (ed.), *Formal Semantics of Natural Language*. Cambridge: Cambridge University Press.

Kamp, H. 1981. ‘The paradox of the heap’. In U. Mönnich (ed.), *Aspects of Philosophical Logic*. Dordrecht: Reidel.

Keefe, R. 1995. ‘Contingent identity and vague identity’. *Analysis* 55: 183-90.

Keefe, R. 2000. *Theories of Vagueness*. Cambridge: Cambridge University Press.

Keefe, R., and Smith, P. (eds.) 1996. *Vagueness: A Reader*. Cambridge MA: MIT Press.

Lewis, D. 1988. ‘Vague identity: Evans misunderstood’. *Analysis* 48: 128-30.

Lewis, D. 1993. ‘Many, but almost one’. In J. Bacon, K. Campbell and L. Reinhardt (eds.), *Ontology, Causality and Mind: Essays in Honour of D. M. Armstrong*. Cambridge: Cambridge University Press.

Lowe, E. J. 1994. ‘Vague identity and quantum indeterminacy’. *Analysis* 54: 110-14.

Lowe, E. J. 1995. ‘The problem of the many and the vagueness of constitution’. *Analysis* 55: 179-82.

Lowe, E. J. 1997. ‘Reply to Noonan on vague identity’. *Analysis* 57: 88-91.

McGee, V., and McLaughlin, B. 1995. ‘Distinctions without a difference’. *Southern Journal of Philosophy* 33: 203-51.

McLaughlin, B. 1997. ‘Supervenience, vagueness and determination’. *Philosophical Perspectives* 11: 209-30.

Noonan, H. 1982. ‘Vague objects’. *Analysis* 42: 3-6.

Noonan, H. 1984. ‘Indefinite identity: a reply to Broome’. *Analysis* 44: 117-21.

Noonan, H. 1990. ‘Vague identity yet again’. *Analysis* 50: 157-62.

Noonan, H. 1995. ‘E. J. Lowe on vague identity and quantum indeterminacy’. *Analysis* 55:

14-19.

Over, D. 1984. ‘The consequences of direct reference’. *Philosophical Books* 25: 1-7.

Over, D. 1989. ‘Vague objects and identity’. *Analysis* 49: 97-9.

Parsons, T. 1987. ‘Entities without identity’. *Philosophical Perspectives* 1: 1-19.

Parsons, T. 2000. *Indeterminate Identity: Metaphysics and Semantics*. Oxford: Clarendon Press.

Parsons, T., and Woodruff, P. 1995. ‘Worldly indeterminacy of identity’. *Proceedings of the Aristotelian Society* 95: 171-91.

Peacocke, C. 1981. ‘Are vague predicates incoherent?’. *Synthese* 46: 121-41.

Pelletier, F. 1984. ‘The not-so-strange modal logic of indeterminacy’. *Logique et Analyse* 27: 415-22.

Pelletier, F. 1989. ‘Another argument against vague objects’. *Journal of Philosophy* 86: 481-92.

Priest, G. 1991. ‘Sorites and identity’. *Logique et Analyse* 34: 293-6.

Priest, G. 1998. ‘Fuzzy identity and local validity’. *Monist* 81: 331-42.

Quine, W. V. 1981. ‘What price bivalence?’. *Journal of Philosophy* 78: 90-5.

Rasmussen, S. 1986. ‘Vague identity’. *Mind* 95: 81-91.

Rolf, B. 1980. ‘A theory of vagueness’. *Journal of Philosophical Logic* 9: 315-25.

Sainsbury, R. M. 1989. ‘What is a vague object?’. *Analysis* 49: 99-103.

Sainsbury, R. M. 1995. ‘Why the world cannot be vague’. *Southern Journal of Philosophy* 33, supplement: 63-81.

Salmon, N. 1982. *Reference and Essence*. Oxford: Blackwell.

Salmon, N. 1984. ‘Fregean theory and the Four Worlds Paradox: a reply to David Over’. *Philosophical Books* 25: 7-11.

Salmon, N. 1986. ‘Modal paradox: parts and counterparts, points and counterpoints’. *Midwest Studies in Philosophy* 11: 75-120.

Sanford, D. 1979. ‘Nostalgia for the ordinary: comments on papers by Unger and Wheeler’. *Synthese 41: 175-84.*

Sanford, D. 1993. ‘The problem of the many, many composition questions, and naive mereology’. *Noûs* 27: 219-28.

Sorensen, R. 1998. ‘Sharp boundaries for blobs’. *Philosophical Studies* 91: 275-95.

Stalnaker, R. 1988. ‘Vague identity’. In D. Austin (ed.), *Philosophical Analysis: A Defense by Example*. Dordrecht: Kluwer.

Thomason, R. 1982. ‘Identity and vagueness’. *Philosophical Studies* 42: 329-32.

Tye, M. 1990. ‘Vague objects’. *Mind* 99: 535-57.

Tye, M. 1994. ‘Sorites paradoxes and the semantics of vagueness’. *Philosophical Perspectives* 8: 189-206.

Tye, M. 1995. ‘Vagueness: welcome to the quicksand’. *Southern Journal of Philosophy* 33, supplement: 1-22.

Tye, M. 1996. ‘Fuzzy realism and the problem of the many’. *Philosophical Studies* 81: 215-25.

Unger, P. 1979a. ‘I do not exist’. In G. Macdonald (ed.), *Perception and Identity*. London: Macmillan.

Unger, P. 1979b. ‘There are no ordinary things’. *Synthese* 41: 117-54.

Unger, P. 1979c. ‘Why there are no people’. *Midwest Studies in Philosophy* 4: 177-222.

Unger, P. 1980. ‘The problem of the many’. *Midwest Studies in Philosophy* 5: 411-67.

Unger, P. 1990. *Identity, Consciousness and Value*. Oxford: Oxford University Press.

van Inwagen, P. 1988. ‘How to reason about vague objects’. *Philosophical Topics* 16: 255-

84.

van Inwagen, P. 1990. *Material Beings*. Ithaca NY: Cornell University Press.

Wheeler, S. 1975. 'Reference and vagueness'. *Synthese* 30: 367-79.

Wheeler, S. 1979. 'On that which is not'. *Synthese* 41: 155-73.

Wiggins, D. 1986. 'On singling out an object determinately'. In P. Pettit and J. McDowell (eds.), *Subject, Thought and Context*. Oxford: Oxford University Press.

Williamson, T. 1994. *Vagueness*. London: Routledge.

Williamson, T. 1995. 'Definiteness and knowability'. *Southern Journal of Philosophy* 33, supplement: 171-91.

Williamson, T. 1996a. 'The necessity and determinacy of distinctness'. In S. Lovibond and S. Williams (eds.), *Essays for David Wiggins: Identity, Truth and Value*. Oxford: Blackwell.

Williamson, T. 1996b. 'What makes it a heap?'. *Erkenntnis* 44: 327-39.

Williamson, T. 1997. 'Imagination, stipulation and vagueness'. *Philosophical Issues* 8: 215-28.

Williamson, T. 1999a. 'On the structure of higher-order vagueness'. *Mind* 108: 127-43.

Williamson, T. 1999b. 'Truthmakers and the converse Barcan formula'. *Dialectica* 53: 253-70.

Williamson, T. 2002. 'Vagueness, identity and Leibniz's Law'. In P. Giaretta, A. Bottani and M. Carrara (eds.), *Individuals, Essence, and Identity: Themes of Analytic Metaphysics*. Dordrecht: Kluwer.

Zemach, E. 1991. 'Vague objects'. *Noûs* 25: 323-40.