

VAGUENESS.

For: Routledge companion to Philosophy of Language, ed. D. Fara and G. Russell.

Abstract

Taking away grains from a heap of rice, at what point is there no longer a heap? It seems small changes – removing a single grain – can't make a difference to whether or not something is a heap; but big changes obviously do. How can this be, since big changes are nothing but small changes chained together?

This article begins by laying out the challenges posed by vague language, and the use of vague language to theorize about vague language. We discuss *epistemicism*—the idea that the phenomena of vagueness are primarily a matter of our *ignorance* of the sharp boundaries of the terms we use. And we sample two *revisionary* theories of vagueness—theories on which the “classical package” of logic and model theory that dominated the twentieth century needs to be tweaked or overhauled altogether to deal adequately with vagueness. Finally, we look more broadly to consider the different tasks a theory of vagueness might address, and how the theories discussed here fit in.

PUZZLES OF VAGUENESS

Taking away grains from a heap of rice, at what point is there no longer a heap? It seems small changes – removing a single grain – can't make a difference to whether or not something is a heap; but big changes obviously do. How can this be, since the big changes are nothing but the small ones chained together? Call this the “little by little” puzzle.

At each stage, while removing grains from the original heap, ask yourself: “is what I have at this moment a heap?”. At the initial stages, the answer will clearly be “yes”. At the late stages, the answer will clearly be “no”. But at intermediate stages, the question generates perplexity: it's not *clearly* right to say “yes”, nor is it *clearly* right to say “no”. A hedged response seems better: “it's a borderline case of a heap”—but that's not yet a *direct* answer. So what is the answer to the original question, when confronted with what we can all agree to be a borderline case of a heap? Call this the “borderlineness” puzzle.

The little by little puzzle leads to the *sorites paradox* (from “sorites” – the Greek word for “heap”). Here's one version. Take a line of 10001 adjacent men, the first with no hairs, the last with 10000 hairs, with each successive man differing from the previous by the addition of a single hair (call this a *sorites series* for “bald”). “Man N” names the man with N hairs.

Man 0 is bald. Man 10000 is not bald. Now consider the following collection of horrible-sounding claims:

- (1): Man 0 is bald, and man 1 is not bald.
- (2): Man 1 is bald, and man 2 is not bald.
- ...
- (10000): Man 99999 is bald, and man 10000 is not bald.

If we *reject* the above, surely we must *accept* their negations:

- (1*): it is not the case that: Man 0 is bald, and man 1 is not bald.
- (2*): it is not the case that: Man 1 is bald, and man 2 is not bald.
- ...
- (10000*): it is not the case that: Man 99999 is bald, and man 10000 is not bald.

But given the various (N*), and the two obvious truths about the extreme cases, a contradiction follows. Each (N*) is (classically) equivalent to the material conditional reading of:

(N**) if Man N-1 is bald, then Man N is bald

Since Man 0 is bald, iterated *modus ponens* reasoning leads to the conclusion that Man 10000 is bald, contrary to our assumptions. QED. We must either find some way of rejecting seemingly compelling premises, or find a flaw in the seemingly valid reasoning.

We turn next to the puzzle of borderlineness: if Harry is intermediate between clear cases and clear non-cases of baldness, “Is Harry bald?” seems to have no good, direct, answer.

There are familiar cases where we cannot answer such questions: having never seen Jimmy I am

in no position to say whether he's bald, simply because I lack relevant knowledge. Lack of knowledge of the answer might also be appealed to in borderline cases. But for such borderline cases, appeal to ignorance simply moves the challenge one stage back. *Why* would we lack knowledge over whether Harry is bald? After all, it seems puzzlement persists even if we find out all the relevant underlying facts about the number and distribution of hairs on Harry's head. What kind of status is this, where the question of baldness is not only something we can't answer, but where we can't even conceive of *getting in a position to answer*?

It's common to say that it's *indeterminate*, or *indefinite*, or *that there's no fact of the matter* whether borderline-bald Harry is bald. Such terminology connects borderlineness to a broader class of putative examples of indeterminacy. Examples include partially defined terms, identities across theory change in science, the problem of the many, ungrounded sentences that generate semantic paradoxes, and indeterminate conditionals. It is a vexed question whether a unified analysis can be given of some or all of these "indeterminacy" phenomena, and what shape it should take if so. (Note too that while a general theory of indeterminacy might answer some questions about the status of borderline cases, it's not at all clear whether it would say anything interesting about the little-by-little puzzle).

SHARP EXTENSIONS

Sometimes it is suggested that vague language poses obstacles to the very *idea* of providing a classical (model-theoretic) semantics for natural language. Is this right?

Focus, for simplicity's sake, on the basic case of extensional semantics. This assigns *extensions* (sets of objects) as the semantic values of predicates. This may be thought to be problematic when the predicates are vague. Sets of objects, as traditionally conceived, are definite totalities—each object is either definitely a member, or definitely not a member. Wouldn't associating one of these totalities with "bald" force us, unjustifiably, to "draw sharp boundaries" for its application? On the other hand, it seems that we *easily* say which set should be the extension of "is bald". It should be $\{x: x \text{ is bald}\}$, i.e. the set of things which are bald. We can *use* vague language ("in the theorist's metalanguage") to *say* what the extension should be.

Delicate questions about methodology in semantics. Can semantic theory say simply that the extension of "red" (in English, or "rouge" in French) is *the set of red things*? Or should we require it give a more illuminating characterization of that set? Even on the latter view, so long as the illuminating characterization is given in vague terms, there's no obvious obstacle to specifying the extension of a vague term. Thus Lewis, in discussing his semantic treatment of *counterfactuals* in terms of the admittedly vague notion of similarity, says "I ... seek to rest an unfixed distinction upon a swaying foundation, claiming that the two sway together rather than independently." (Lewis, 1973, p.92). It would be problematic if one imposed the requirement that we should be able to give a semantics for a vague object language within a *non-vague* metalanguage. **But** such a requirement would need to be argued for.

Perhaps, though, the worry is that terms such as $\{x: x \text{ is red}\}$ picks out a "vague set" – an entity not covered by standard set theories. Consider, for example, the set of bald men B, and the set B+ that is the union of B with {Harry}, where Harry is borderline bald. Since it's indefinite whether Harry is bald, it seems it will be indefinite whether Harry is a member of B; but he's

definitely a member of B^+ . All other individuals will be a member of B iff they are a member of B^+ . It follows that it is indefinite whether B and B^+ contain exactly the same members. By the axiom of extensionality, two sets are identical iff they have the same members. So it seems as if it will be indefinite whether B and B^+ are identical.

This may be problematic. Indefiniteness in identity is the target of the famous Evans argument (Evans 1978, Salmon 1981). This purports to show, via an appeal to Leibniz's law, that such claims can be reduced to absurdity: facts about identity are never vague. So, not only do we owe an account of vague sets in general, we also need find some way of addressing Evans-Salmon argument.

However, a friend of textbook (classical) semantics has responses. One idea is that terms such as “ $\{x: x \text{ is red}\}$ ” are *vague* descriptions of *precise* sets. Suppose the relevant domain D contains a single object a , and it is indefinite whether a is red. The candidates to be picked out by “ $\{x \text{ in } D: x \text{ is red}\}$ ” are the null set and the singleton $\{a\}$. It might be indefinite which of the two it picks out—but a clear mistake to think it picks out some special entity, a “vague set”, indefinitely identical to both. Indeterminacy in reference is a well-known escape route from the Evans reductio argument (cf. Lewis 1988). If this line succeeds, no retreat from classical conception of sharply-bounded sets would be needed in order to appeal to underpin their use in semantics.

However, even if textbook semantics were fine, it wouldn't *explain* the two puzzles of the previous section: borderlineness and the sorites. In the next three sections I present examples of three kinds of positive accounts of vagueness and how they tackle the puzzles identified above. I will start with the one that requires least modification of standard logic and semantics – epistemicism. Each is successively more revisionary.

CLASSICAL VAGUENESS: EPISTEMICISM

If we stick with classical logic and semantics for vague language, a consequence is that room for maneuver in explaining the puzzles of borderlineness and the sorites paradox is drastically reduced. Classicists cannot, for example, say that borderline-bald Harry is special because it is neither true nor false that he is bald – for this would conflict with the classical commitment to *bivalence* (every claim being either true or false). And unless we want to deny that Jerry Garcia in his prime was bald, and Kojak was not, then it follows, classically, that in the baldness sorites series, the last non-bald man differs by only a hair from a man who is bald.

The last claim seems particularly unpalatable – it seems to commit us, incredibly, to sharp boundaries to baldness. Some endorsing the classical logic/semantics package argue that no such commitment follows (McGee & McLaughlin 1994); but the theorists I discuss in this section take the hard-nosed stance that there really are sharp boundaries for vague predicates. To make a case for this surprising claim, at a minimum we need some explanation of how the characteristic features of vagueness are consistent with sharp boundaries.

Epistemicists typically endorse the face-value classicism just described. It's either true, or false, that borderline-bald Harry is bald, for example – there is a fact of the matter – but we do not and (in a certain sense) cannot know which option is realized. Below, I sketch Williamson's (1994) elaboration of this idea.

Let us consider borderlineness first. Start from the idea that we are ignorant of whether Harry is bald, when he is a borderline case. The puzzle was to explain why this was so, and why the unknowability was of such a strong and ineliminable sort.

Williamson's proposal makes use of a general constraint on knowledge: the idea that in order to know that *p*, it cannot be a matter of luck that one's belief that *p* is true. Williamson articulates this as a "safety principle", roughly:

(SAFETY)

For "S knows that *p*" to be true (in such situation *s*), "*p*" must be true in any marginally different situation *s** (where one forms the same beliefs using the same methods) in which "S believes *p*" is true.

The idea is that the situations *s** represent "easy possibilities": falsity at an easy possibility makes a true belief too lucky to count as knowledge.

This first element of Williamson's view is independently motivated epistemology. The second element is that the extensions of vague predicates, though *sharp*, are *unstable*. They depend on exact details of the patterns of use of vague predicates, and small shifts in the latter can induce small shifts in the (sharp) boundaries of vague predicates.

Given these two, we can explain our ignorance in borderline cases. A borderline case of "bald" is just one where the boundary of "bald" is close enough that a marginally different pattern of usage could induce a switch from (say) Harry being a member of the extension of "bald" to his not being in the extension of that predicate. If that's the case, then even if one truly believed that Harry was bald, there will be an easy possibility where one forms the same beliefs for the same reasons, but that sentence is false. Applying (SAFETY), the belief won't count as knowledge.

Given that the source of ignorance resides in the sharp but unstable boundaries of vague predicates, one can see why gathering information about hair-distributions won't overcome the relevant obstacle to knowledge. This is why the ignorance in borderline cases seems ineliminable.

What does the epistemicist say about the sorites? Whether we present that argument via a string of negated conjunctions, or conditionals, or variations and generalizations thereof, one of the premises will be false: a certain man will be bald, while his neighbour with just one more hair is not bald. The sorites argument is therefore unsound. But it's controversial whether this is enough to *resolve* our initial little-by-little puzzle. We'd like to know why we found the idea of a sharp cut off so incredible (especially since there's a very simple, valid argument from obvious premises to this effect available). Williamson can give an account of why we'd never feel able to accept – since we can never know – any one of the individual conjunctions. But that doesn't explain why we're uneasy (to say the least) with the thought that *some such conjunction is true*. An analogy: I'll never know in advance which ticket will win a lottery; but I'm entirely comfortable with the thought that *one will win*. Why don't we feel the same about the sorites?

SUPERVALUATIONISM

A common reaction to borderline-bald Harry is that it's *neither true nor false* that he is bald.

Assuming that one can only know what is true, this would explain our inevitable lack of knowledge in borderline cases. It's often thought to be a rather plausible suggestion in itself.

Classical semantics builds in the principle that each meaningful claim is either true or false. So if we're to pursue the thought that borderline claims are truth value gaps, we must revise our semantic framework to some extent. Indeed, we can know in advance that any semantic theory with truth-value gaps will diverge from classical semantics even on some of the most intuitively plausible consequences: for it can be shown under very weak assumptions that truth value gaps are incompatible with accepting disquotational principles such as: "Harry is bald" is true if and only if Harry is bald (see Williamson 1994, ch7).

How will the alteration of the classical framework go? One suggestion goes under the heading "supervaluationism" (see *inter alia* Fine 1975, Keefe 2000). As we'll see, the term is somewhat ambiguous.

As an account of the *nature* of vagueness, supervaluationism is a view on which borderlineness arises from what we might call "semantic indecision" (cf. Lewis 1993). Think of the sort of things that might fix the meanings of words: conventions to apply the word "bald" to clear cases; conventions to apply "not bald" to clear non-cases; various conventions of a more complex sort—for example, that anyone with less hair than a bald person should count as bald. The idea is that when we list these and other principles constraining correct interpretation, we'll be able to narrow down the space of acceptable (and entirely classical) interpretations of English—but not to the single *intended* interpretation hypothesized by classical semantics. At best, what we'll get is a cluster of candidates. Let's call these *the sharpenings* for English (sometimes these are called "precisifications" or "delineations"). Each will assign to each vague predicate a sharp boundary. But very plausibly the location of such a boundary is something the different sharpenings will disagree about. A sentence is indeterminate (and if it involves a vague predicate, is a borderline case) just in case there's a sharpening on which it comes out true, and another on which it comes out false.

As an account of the *semantics* of vague language, the core of the supervaluationist proposal is a generalization of the idea found in classical semantics, that for something to be true is for it to be true *at the intended interpretation*. Supervaluationism offers a replacement. It works with a set of "co-intended interpretations", and says that for a sentence to be true, it must be true at *all the co-intended interpretations* (this is sometimes called "supertruth"). This dovetails nicely with the semantic indecision picture, since we can take the "co-intended interpretations" to be what we called above the *sharpenings*. When a sentence is indeterminate (true on one sharpening and false on another) neither it nor its negation will be true: and hence we have a *truth value gap*.

The slight tweak to the classical picture leaves a lot unchanged. Consider the tautologies of classical logic, for example. Every *classical* interpretation will make them true; and so each sharpening is guaranteed to make them true. Hence, any classical tautology is always supertrue, which is enough to make it a *supervaluational* tautology. (It's a matter of dispute whether more subtle departures from classical logic are required, and whether this matters. Cf. (Fine 1975, Williamson 1994 ch 5, Williams 2008)).

If (super)truth is a constraint on knowledge, supervaluationists can explain why we can't know whether borderline bald Harry is bald. On some developments of the position, they can go

interestingly beyond this. One might argue that insofar as one should only invest credence in a claim to the extent one believes it true, obvious truth-value-gaps are cases where we should *utterly* reject (invest no credence in) both the claim and its negation (cf. Field 2003). This would mean the information that such-and-such is borderline gives us a direct fix on what our *degree of belief* should be in borderline cases.

What about the sorites? No individual conjunction (N) “Man N is bald and man N+1 is not bald” will be true—many of them will instead be *truth value gaps*, true on some sharpenings and false on others. On the other hand, *every* sharpening can provide a number n that makes “for some n, man n is bald and man n+1 is not bald” true—different ones on different sharpenings—so it’s supertrue overall. This highlights one of the distinctive (and perhaps disturbing) features of supervaluationism—that disjunctions and existential generalizations can be true, even if no disjunct or instance is. Either way, not all the premises of the argument for paradox will be true, so the argument is blocked (just as for the epistemicist, there is a *further* question about what accounts for the argument’s original appeal).

Two points to bear in mind about supervaluationism. First, as we’ve seen, the supervaluationist *endorses* the claim that “there is a cut-off for bald” – a pair of men differing by only one hair, with the first bald and the second not. The supervaluationist must try to persuade you that once one understands the sense in which *there is no fact of the matter* where that cut-off is, you will be prepared to accept the existence of cut-offs. Second, many want to press the charge that the supervaluationist makes no progress over the classicist, for reasons of “higher order vagueness”. The thought is that the task of explaining how a *set of sharpenings* gets selected by the meaning-fixing-facts is no easier or harder than explaining how a *single classical interpretation* gets picked out. However, (a) the supervaluationist can reasonably argue that if she spells out the notion of “sharpening” in a vague metalanguage, she will regard the boundary between the sharpenings and non-sharpenings as vague (see Keefe 2000); (b) *even if* both epistemicist and supervaluationist were both in some sense “committed to sharp boundaries”, the account they give of the nature of vagueness is vastly different, and we can evaluate their positive views on their own merits.

MANY-VALUED SETTINGS

A distinctive feature of supervaluationism was that while it threw out *bivalence* (“Harry is bald” is either true or false) it preserved the corresponding instance of *excluded middle* (“Harry is bald or Harry is not bald”). Revising logic and semantics in a more thorough-going way would allow for a coherent picture where we can finally reject the claim “there is a single hair that makes the difference between bald and non-bald” without falling into paradox.

“Many valued” logics can be characterized by *increasing the number of truth-values we work with*—perhaps to three, perhaps infinitely many—and offering generalizations of the familiar stories of how logical constants behave to accommodate this tweak. There are many ways to develop this (for a survey, see *Many Valued Semantics*, this volume).

Here is a sample many-valued logic, for a propositional language with conjunctions, disjunctions and negations. To characterize the logic, we postulate three values: call them, neutrally, “1” “0.5” and “0”. For the propositional case, each atomic sentence will be assigned one of these

truth values. The truth values then get assigned to complex sentences recursively. Let us assume a conjunction will be assigned the minimum of the truth values of its conjuncts; a disjunction will get the maximum of the truth values of its disjuncts; and a negation will be 1 minus the truth value of the claim negated (you can easily check that, ignoring the value 0.5 altogether, we get back exactly classical truth-tables.)

One option at this point is to *take this model theory seriously*—much as the classicist and supervaluationist do—and hypothesise that natural language has (or is modelled by) some many-valued interpretation (or set of interpretations). The most famous version of this many-valued approach, the “degree theory” of vagueness, works not with the three values above, but *infinitely many*, labelled by the real numbers between 1 and 0. The truth values are thought of as “degrees of truth” – with a sentence getting closer to “perfect truth” insofar as its truth value approaches 1. (Cf. inter alia Machina 1976, Smith 2008)

Degree theorists will say that a borderline case of a vague predicate has an intermediate truth value. *Claiming that Harry is bald* is neither wholly true, nor wholly false—and this is the reason why we are uncomfortable giving either answer. Along a sorites series, the degree of truth of “Man N is bald” falls off smoothly: for N=0 this may be value 1; for large N is value close to 0. What one says about the sorites argument itself—the truth values of the compound claims involved, and the validity of the argument, will depend on exactly what story is given about the behaviour of the connectives and the characterization of the logic.

Worries about “higher order” vagueness are more pressing directed against the degree theorist than against the supervaluationist. On the degree-theoretic story, it appears that Harry will be bald to some particular, precise degree – say 0.5456. But is it really plausible that there are such sharp facts of the matter about the *exact degree* to which Harry is bald? The “higher order indeterminacy” claim is that it is *indeterminate* to what degree Harry is bald. But if this higher-order indeterminacy is handled in a many-valued manner, then the language in which we ascribe degrees of truth to sentences – and hence theorize about degree theory itself – is non-classical. It’s then not clear that the kind of claims we made earlier to characterize the position, e.g. “every atomic sentence has a degree of truth”, are fully true. Many-valued model theory is usually described in a metalanguage tacitly assumed to be classical. It’s far from clear this is dispensable (cf. Williamson 1994 ch.4).

Some claim many-valued theories give counterintuitive results. Take a 0.5-valued sentence A. From the recursive clauses, we can see that “A&~A” will be half-true as well. Many find this commitment objectionable. (See Williamson 1994 ch.4 for arguments that *all* many-valued treatments of the connectives will be similarly problematic). Degree-theorists don’t agree, and stand-off threatens. It’s worth noting that this objection relies on a particular, “degrees of truth” gloss on the many-valued semantics. But in a three-valued setting, rather than talk of degrees of truth, one could interpret value 1 as “true”, 0 as “false” and 0.5 as “neither true nor false” (cf. Parsons 2000). On this reading, “A&~A” will be described as a truth value gap, when “A” itself has this status – which sounds rather different (and more plausible) than describing it as half-true. The opponent of many-valued approaches must check that their objections don’t depend on a contentious gloss on the targeted view.

We've concentrated till now on accounts that *take model theory seriously* – and which look for an integrated account of the logic and semantics of vague language. But there's another option – to focus attention on the logic, and regard model-theory as simply a descriptive instrument we use to specify the One True Logic.

Logic-first theorists advocate a revisionary take on what *arguments are valid* and what *combinations of claims or attitudes are consistent*. To fix ideas let's focus on an example, using the 3-valued assignments described earlier. Recall that a sequent is valid on the *strong Kleene logic* if it is "1-preserving" across these assignments, i.e. if it is such that when all the premises are value 1, the conclusion is value 1 too. In this Kleene logic, " $A \vee \sim A$ " is no longer a tautology, since it need not be value 1. On the other hand, " $A \& \sim A$ " is still treated as a logical contradiction (every sentence whatsoever follows from it), since it will never attain value 1, no matter what value A is assigned.

For the Assignment-first theorists we have been considering thus far, the interesting question was about the various assignments and what they represent. Does the middle status represent a truth value gap or a degree of truth? Among the variety of assignments compatible with the rules, which matches up to the actual distribution of truth statuses? But for a logic-first theorist, such questions are misguided. The assignments are simply mappings from sentences to numbers – there's no need to think that they latch on to real properties of the sentences, and so no sense in asking about what they represent or which one among them is "intended" – all they do is collectively characterized a (Kleene-style) logic.

A logic-first theorist such as Field (2003a) then uses the logic to handle the distinctive puzzles of vagueness. Field is committed to a certain *normative significance* for logic – that accepting the premise of a valid argument commits one to accepting its conclusion; and rejecting the conclusion commits one to rejecting its premise. Note that the Kleene-logic does not make excluded middle a tautology, so on Field's view there's no *logical* reason to accept instances of *it*. Field advocates *rejecting* the claim that "Harry is bald or he isn't" when Harry is borderline-bald. Since this disjunction follows, in the logic he endorses, from "Harry is bald", he rejects this claim too, as well as its negation. This then is Field's account of the puzzle of borderlineness: both direct answers to "Is Harry bald?" are to be rejected.

Note that this story would collapse if to reject a claim is simply to accept its negation – we must be able to reject "Harry is bald" without accepting "Harry is not bald". That situation can get more dramatic. In the Kleene logic, the negation of excluded middle " $\sim(A \vee \sim A)$ " is inconsistent (and so must never be accepted). But we are supposed to reject instances of $A \vee \sim A$. So a basic feature of the position is that sometimes we reject claims whose negations are inconsistent.

One can view the sorites paradox as repeating this pattern on a grander scale. Instances of "Man N is bald and Man N+1 isn't" do indeed sound horrible, and we should reject them. But the sorites paradox does not get started unless we move from this to endorsing their negations. What the sorites paradox reveals is that the negations generate a contradiction. But just as with a single instance of excluded middle, there's no obligation to move from rejecting the claim to endorsing its (inconsistent) negation.

Clearly, the logic-first version of many-valuedism is very different from the assignment-first

version. Note, for example, that the higher-order indeterminacy worries mentioned earlier simply do not get started against the logic-first theorists (at least if they endorse the instrumental view of model theory). If one's account has no place for "the intended truth value assignment", then opponents can hardly ask embarrassing questions about whether it's determinate which assignment this is!

On the other hand, the logic-first position is radical. It gives up on the attractive prospect of a model-theoretic *analysis* and *explanation* of validity. And it seems committed to doing without model theory when directly addressing questions about the semantics and truth status of vague language (Field, for example, endorses a non-semantic, *disquotational* treatment of truth and meaning).

CONCLUSION

I've outlined three sample positive views of vagueness. They are illustrative of broad *categories* of theories of vagueness. The categories are not exhaustive, but they do orientate much contemporary debate.

The most conservative category—represented here by epistemicism—preserves classical logic, model theory and textbook semantics. The distinctive *epistemicist* response is one way of explaining what vagueness could be, even if the apparent commitments of this setting to "sharp cut-offs" are taken at face value. But this isn't the only way a classicist might respond. Fara (2003), for example, agrees with Williamson on the *broad classical setting*. But her account of what vagueness is – at least for the cases she focuses on, gradable adjectives – is very different. Her idea is that (i) the word "red" means something like *significantly redder than is typical*; (ii) what counts as "significant" is deeply interest-relative, and in particular, can vary depending on what is the focus of our attention; (iii) as a result of this, we can predict that whenever we focus on a particular potential "cut-off" for red, the facts about what is now *significant* to us ensure that the classical cut-off for "red" is located somewhere *other* than we are looking. Hence, for Fara, the distinctive elusiveness of vague predicates. (The idea that context-sensitivity might be intimately involved with the phenomena of vagueness is something many writers find appealing, whether or not they work in a classical backdrop (see Raffman 1994, Shapiro 2006)).

The next most conservative response is to design a "special purpose" semantics for vagueness, allowing semantic features to do significant explanatory work, which nevertheless "saves much of the appearance" of the standard classical setting. Supervaluationism illustrates this possibility. But there are other ways to go, even with the broad "semantic indecision" framework that gives us our notion of a *sharpening*. Various authors have suggested that we might look at the *proportion* of sharpenings on which a given claim holds. In terms of this, we can characterize a notion of *intermediate degrees of truth*: the degree of truth of "Harry is bald" being the proportion of sharpenings on which it is true. Thus supertrue sentences ("Kojak is bald") will be true to degree 1; while if "Harry is bald" is plumb borderline, true on half and false on half, then it'll be true to degree 0.5 (see Edgington 1997). Such notions may find application within a treatment of comparative forms of vague adjectives and modifiers (see Kamp 1975; Lewis 1970).

The final theoretical framework we looked at was thoroughly revisionary, though the revisionism

has more or less radical forms. Preserving realism about model-theoretic foundations, we find “gap” and “degree” theories. Alternatively, looking at model-theory as a mere descriptive instrument for specifying a logic, we have logic-first theories. There’s plenty of room for debate about which logic is most appropriate: for example, Crispin Wright (2003, 2007) argues for *intuitionism* as the appropriate logic governing vague discourse. (Wright has been a long-time advocate of drawing lessons on the nature of semantic theory from careful study of the phenomenon of vagueness—see Wright (2007)).

Important though these logico-semantic issues are, one should not think that the whole task of understanding vagueness consists in figuring out the proper semantic treatment of it. As is evident from the case of epistemicism, resources outside the philosophy of language may be key to diagnosing its characteristic puzzles. Moreover, logico-semantics may *underdetermine* the overall theory of vagueness. Getting clear on the *nature* of vagueness/indeterminacy is another aspect to theorizing about vagueness—for example is it a matter of “semantic indecision” as many supervaluationists contend? Similar formal machinery can be paired with very different responses to this question. As well as the *nature, psychological* features demands attention. Should we *have zero credence* in Harry being bald, if we know he’s a borderline case (as Field argues)? Or should have fifty-fifty confidence, as some degree theorists maintain? (Cf. Smith 2008, Schiffer 2003.) Or something else entirely (Wright 2003)?

The attention lavished on vagueness since the mid-1970’s shows little signs of converging on a single standard account. But it has given us plenty of options, providing rich resources for future arguments. To study it is to appreciate the far-reaching impact of the disarmingly simple puzzles of vagueness.

Bibliography

D. Edgington. (1997) "Vagueness by Degrees," in Vagueness: A reader, ed. Rosanna Keefe and Peter Smith (Cambridge, MA: MIT Press), 294-316.

G. Evans. (1978) "Can there be vague objects?" Analysis 38, no. 13: 208.

D. Graff Fara. (2003) "Shifting Sands: an interest-relative theory of vagueness," Philosophical Topics 28, no. 1: 45-81.

H. H. Field. (2003) "Semantic paradoxes and the paradoxes of vagueness," in Liars and Heaps, ed. J. C Beall (2003) 262-311.

H. H. Field. (2003a) "No fact of the matter," Australasian Journal of Philosophy 81, no. 4: 457-480.

K. Fine. (1975) "Vagueness, truth and logic," Synthese 30: 265-300.

J. A. W. Kamp. (1975) "Two theories about adjectives," in Formal Semantics of Natural Language, ed. E. Keenan (Cambridge: Cambridge University Press), 123-155.

R. Keefe. (2000) Theories of Vagueness (Cambridge: Cambridge University Press).

D. K. Lewis. (1970) "General Semantics," Synthese 22: 18-67.

--- (1973) Counterfactuals (Oxford: Blackwell).

--- (1988) "'Vague Identity: Evans misunderstood'," Analysis 48: 128-30.

--- (1993) "Many, but almost one," in Ontology, Causality and Mind: Essays on the philosophy of D. M. Armstrong, ed. Keith Campbell, John Bacon, and Lloyd Reinhardt (Cambridge: Cambridge University Press).

K. F. Machina. (1976) "Truth, belief and vagueness," Journal of Philosophical Logic 5: 47-78.

Vann McGee and Brian McLaughlin. (1994) "Distinctions without a difference," Southern Journal of Philosophy supp XXXII 203-251.

D. Raffman. (1994) "Vagueness without Paradox," Philosophical Review 103.

N. U Salmon. (1981) Reference and essence (Princeton University Press).

S. Schiffer. (2003) The Things We Mean (Oxford University Press, US).

S. Shapiro. (2006) Vagueness in context (Oxford University Press, USA).

N. J. J. Smith. (2008) Vagueness and degrees of truth (Oxford: Oxford University Press).

J. R. G. Williams. (2008) "Supervaluationism and logical revisionism," The Journal of

Philosophy CV, no. 4.

T. Williamson. (1994) Vagueness (London: Routledge).

C. Wright. (2003) “Vagueness: A fifth column approach,” in Liars and Heaps: New essays on paradox, ed. J. C. Beall (Oxford: Oxford University Press), 84-105.

--- (2007) “Wang’s Paradox”, in The Philosophy of Michael Dummett, ed. Auxier, Hahn, (Open Court).