



Imagination, Stipulation and Vagueness

Timothy Williamson

Philosophical Issues, Vol. 8, Truth. (1997), pp. 215-228.

Stable URL:

<http://links.jstor.org/sici?sici=1533-6077%281997%298%3C215%3AISAV%3E2.0.CO%3B2-1>

Philosophical Issues is currently published by Ridgeview Publishing Company.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/rpc.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

Imagination, Stipulation and Vagueness

Timothy Williamson

1

Humans are better at logic than at philosophy.¹ When philosophical considerations leads someone to propose a revision of basic logic, the philosophy is more likely to be at fault than the logic. Although there is no general methodological ban on such a revision, just as there is no general methodological ban on a philosophically motivated revision of basic physics, any particular revisionary proposal must be greeted with some scepticism. Even before we have found the fault in the philosophical argument, we should think it likely that there is one—which is not to say that we need not bother to look for it.²

The epistemic view of vagueness is based on the use of classical logic in vague languages, together with disquotational principles for truth and falsity. The latter may be taken as these schemata (where names of sentences replace 's'):

TRUE If *s* says that *P*, then *s* is true if and only if *P*.

¹Humans are *collectively* better at logic than at philosophy. Of course, some individual humans are better at philosophy than at logic.

²Michael Tye plausibly suggests that humans are also better at non-classical logic than at philosophy. Indeed they are, when they use a classical metalogic.

FALSE If s says that P , then s is false if and only if not P .

TRUE and FALSE may be regarded as giving the basic logic of 'true' and 'false' as applied to sentences. For example, if the sentence 'Jack is bald' says that Jack is bald, then 'Jack is bald' is true if and only if Jack is bald, and false if and only if Jack is not bald. By the classical law of excluded middle, either Jack is bald or Jack is not bald. By more classical logic, it follows that 'Jack is bald' is either true or false, and that either 'Jack is bald' is true or 'Jack is not bald' is. Suppose that Jack is a borderline case for 'bald'. Then, uncontentiously, we have no idea how to find out whether Jack is bald—which is not yet to say that there is something to be ignorant of here. But if either 'Jack is bald' is true or its negation is, then there is something to be ignorant of. There is a truth that we cannot know—at least, not in any ordinary way; whether we could know it on the testimony of a *deus ex machina* is not a point of much consequence.

Some try to reconcile bivalence and the undecidability of borderline cases with the denial of 'genuine' ignorance. The problem lies in explaining a clear alternative to the epistemic account of borderline cases. A careful recent attempt by Vann McGee and Brian McLaughlin illustrates the difficulties. They discern a clash in borderline cases between the principle of disquotation and the principle of correspondence, two fundamental components of our conception of truth. Their disquotation principle conjoins homophonic disquotational principles about truth and falsity. Their correspondence principle says

that the truth conditions for a sentence are established by the thoughts and practices of the speakers of the language, and that a sentence is true only if the nonlinguistic facts [their footnote: Assuming the sentence is not about language] determine that these conditions are met. (1995: 214)

According to McGee and McLaughlin, if 'Harry is bald' is a borderline case, then the correspondence principle implies that it cannot be either true or false (215). Their disquotational principle implies that it is either true or false. They suggest that we should retain only the disquotation principle for truth and falsity, but use the correspondence principle to characterize stronger notions of *definite* truth and *definite* falsity (217). Thus 'Harry is bald' is either true or false, but neither definitely true nor definitely false. A supervaluationist semantics is provided for an object-language definiteness operator 'Def', on which [Def P] is true in a [classical] model if and

only if [P] is true in every admissible [classical] model.³ Classical logic is preserved. The correspondence principle appears not to involve epistemic notions, and therefore to be eligible to characterize a non-epistemic notion of definiteness. But consider the supposed conflict between the disquotation and correspondence principles. One can argue as follows:

- (*) The thoughts and practices of speakers of English establish that the truth condition for 'Harry is bald' is that Harry is bald, and that its falsity condition is that Harry is not bald. Either Harry is bald or Harry is not bald. If Harry is bald, then it is a nonlinguistic fact that Harry is bald, and that fact determines that the truth condition for 'Harry is bald' is met. If Harry is not bald, then it is a nonlinguistic fact that Harry is not bald, and that fact determines that the falsity condition for 'Harry is bald' is met. Thus the correspondence principle permits 'Harry is bald' to be either true or false.

McGee and McLaughlin must reject (*). Nevertheless, (*) shows that there is no conflict in borderline cases between the disquotation and correspondence principles as such. McGee and McLaughlin must be assuming that even if Harry is bald, it is not a genuine nonlinguistic fact that Harry is bald, or perhaps that the condition that Harry is bald is not a genuine truth condition. This unexplained implicit standard of genuineness is what gives bite to the correspondence principle and the notion of definiteness. Until the standard is explained, it would be naive to assume that the underlying condition for it to be met is as non-epistemic as the authors intend. Thus McGee and McLaughlin have not succeeded in articulating a clear alternative to the epistemic view of vagueness.⁴

The initial case for the epistemic view is powerful. Nevertheless, many philosophers regard the view as too counterintuitive to be taken seriously. Is epistemicism about vagueness therefore like David Lewis's modal realism and Graham Priest's dialetheism — hard to refute, hard to believe, the victim of the incredulous stare? One crucial difference is that modal realism and dialetheism, unlike epistemicism, are revisionary in logic. Dialetheism says that some

³The combination of such a supervaluationist account of definiteness with a disquotational notion of truth, which is *not* truth in all admissible models, was anticipated by Kit Fine, 1975: 296.

⁴Williamson 1995 discusses the argument from classical logic and semantics to the epistemic view in more detail.

contradictions are true; faced with the trivializing classical deduction of any proposition whatsoever from a contradiction, it rejects classical logic (Priest 1987). Modal realism treats talk about possible identity by means of counterpart theory, and rejects the thesis of quantified modal logic (with a classical logic of identity) that identicals could not have been distinct (Lewis 1986: 263). Epistemicism employs a different methodology: one holds one's logic fixed, to discipline one's philosophical thinking. It is its opponents who reject the discipline. The epistemicist's hunch is that in the long run the results of the discipline will be more satisfying from a philosophical as well as from a logical point of view.

What is the source of the anti-epistemicist intuitions? §II argues that the sharp cut-off points for vague terms implied by the epistemic view are in a sense unimaginable, which makes the view counterintuitive without constituting an argument against it. Some philosophers find the epistemic view particularly counterintuitive when applied to simple cases of incomplete stipulations. §III argues that the view permits a coherent account of such cases. Obviously, these two points constitute nothing like a full case for the epistemic view.⁵ Their purpose is to suggest that we shall misunderstand what our intuitions have to teach us about vagueness if we treat them in an uncritical spirit.

2

Imagine a heap of sand. Imagine one grain removed, the rest undisturbed. What is left is a heap, in your imagination (or so I conjecture). It seems imaginatively compelling that the result of removing one grain from a heap is still a heap.⁶ A sorites paradox then ensues in the usual way. But how strongly does this imaginative exercise support the major premise of the sorites paradox, that the result of removing one grain from a heap is still a heap?

Imagine a bird. It is not a penguin, in your imagination (or so I conjecture). Clearly, that imaginative exercise gives no serious support to the claim that penguins are not birds. They are not *prototypical* birds, they are not the kind of bird which you imagine when you are simply asked to imagine a bird, but nevertheless they quite definitely are birds. What do you do when you are asked to

⁵For a less incomplete case for epistemicism see Williamson 1994.

⁶The assumption that the other grains are undisturbed will henceforth be tacit.

imagine a heap? Very likely, you imagine a prototypical heap, one with more than enough grains to constitute a heap. So the earlier imaginative exercise may not support anything stronger than this claim: the result of removing one grain from a *prototypical* heap is still a heap (not necessarily a prototypical one). That claim generates no sorites paradox; it does not even entail that the result of removing two grains from a prototypical heap is a heap.⁷

It would be premature to conclude that the imagination is useless when we think about vagueness. But its use must be carefully controlled by an analysis of its limitations. In particular, this question arises: if one's imagination balks at the epistemic view of vagueness, is it because the epistemic view is false, or because one's imagination is limited? In what follows, it will be argued that since the epistemic view predicts that the sharp cut-off points which it postulates are unimaginable in a certain sense, the discovery that they are indeed so unimaginable does not constitute evidence against the epistemic view. Furthermore, since there is no reason to suppose that our imaginations can encompass everything that there is, the prior probability of the epistemic view is not diminished by its prediction of unimaginable cut-off points.

In this context, the earlier argument about prototypes is no longer sufficient. The objection to the epistemic view is that our imaginations fail when we try to imagine a heap being turned into a non-heap by the removal of one grain. Since the specification of the task obviously requires the heap to be near the borderline for heaphood, the default choice of a prototypical heap is blocked. Our imaginative failure has a deeper source.

Consider the dismantling of a heap, grain by grain. On the epistemic view, the process contains a stage at which a heap is turned into a non-heap by the removal of one grain. Call any such stage a *transition*. Can we imagine a transition? 'Imagine' has a range of senses. At one extreme, to imagine that P is merely to suppose that P, if only as a hypothesis for *reductio ad absurdum*. In that sense, we can certainly imagine that a given stage in the dismantling process is a transition. But that does nothing to show that it could *be* a transition. Any proposition can be made a hypothesis. If we cannot deduce a contradiction from the hypothesis that the given stage is a transition, that fact contributes to the defence of the epistemic view, but it does not address the sense in which our imaginations *do* balk at the hypothesis.

⁷For a recent exchange on prototypicality and vagueness, see Kamp and Partee 1995, Fodor and LePore 1996 and Osherson and Smith, forthcoming.

At the other extreme, to imagine something is a substitute for experiencing it. To imagine a building from the architect's plans is a substitute for the experience of walking round and looking at it. To imagine feeling angry is a substitute for the experience of feeling angry. Of course, the word 'substitute' here is just a gesture towards a proper account, which I will not attempt to supply. But this much is implied: the cognitive point of the imagination, in this sense, is that by imagining something one can gain some of the information about it that one could have gained by experiencing it. The relevant ways of experiencing something include perceiving it, feeling it and (in an extended sense of 'experience') remembering it. Conversely, by imagining something one cannot expect to gain information about it that one could not in principle have gained by experiencing it. In particular, if one cannot in principle recognize a stage as transitional for heaphood when one experiences the dismantling process, then one cannot expect to recognize it as transitional when one imagines the process. One can still imagine *that* the stage is transitional, but that is only in the suppositional sense of 'imagine' whose irrelevance has already been noted.

On the epistemic view, one cannot recognize a stage as transitional when one experiences the dismantling process. Why not? The answer is not that the transition lacks a distinctive appearance, one different from the appearances of all other stages. The epistemic view cannot allow that two successive stages of the process always have the same appearance, for then the first and last stages would have the same appearance, which they patently have not. The logic of identity makes having the same appearance a transitive relation. Perceptual *indiscriminability* is a non-transitive relation, but it is not the case that in every sorites series any two successive stages of the dismantling process are perceptually indiscriminable. To make the point vivid, let the initial heap consist of twenty large distinctively shaped and coloured stones. Each stage in the process can easily be discriminated from every other by the naked eye. On the epistemic view, one stage is in fact transitional. Whichever it is, that stage is perceptually discriminable from every other stage. One can recognize it as that stage. Nevertheless, the stages form a sorites series; one cannot recognize the transition as transitional, for that would be to discover the cut-off point for heaphood.

The block to recognizing the transition as a transition is conceptual, not perceptual. The full epistemic story is too complex to recount here; some remarks must suffice. Consider a moment just before the transition. The judgement that there is a heap is still true. Someone who makes it correctly applies the concept of a heap. But

the making of the judgement is not perfectly sensitive to the total pattern of dispositions and causal connections with the environment in virtue of which the concept being exercised is that of a heap rather than that of a heap*, where it takes one more grain (or stone) to make a heap* than it does to make a heap. The content of a genuine concept is not exhausted by the fact of its present application. If the pattern of use had been very slightly different, the concept possessed would have been that of a heap* rather than that of a heap, and the person might very easily have exercised that concept in the same situation in the same way, thereby judging that there is a heap*. The latter judgement would have been false, for since there are only just enough grains to make a heap, there are not quite enough to make a heap*. Thus the judgement that there is a heap issues from a method of judgement which could very easily have issued in a false judgement in the same situation.^{8,9} The basis of the true judgement that there is a heap is highly unreliable. Thus the judgement fails to express knowledge; in this situation, one cannot know that there is a heap. The removal of the next grain is the transition, but recognizing it as a transition would involve knowing in the present situation that there is a heap. Thus one cannot recognize the transition as a transition.

Since one cannot recognize a stage as transitional when one *experiences* the dismantling process, one cannot recognize it as transitional when one *imagines* the process. Thus the epistemic view predicts that sharp cut-off points are unimaginable in the relevant sense. That sharp cut-off points are indeed unimaginable is some slight empirical confirmation of the epistemic view. Of course, any independent reason to hold that if sharp cut-off points exist then they are imaginable would constitute a prior objection to the epistemic view, but we have no such reason. Our imagination is not the measure of all things. Since our intuitions depend on our imaginative capacities, there will always be something unintuitive about the epistemic view. But we are adults; once we have traced the unintuitiveness to its source, we can learn to give it no more philosophical weight than it deserves.

The same dialectic applies at the metalinguistic level. Rather than focus on the relation between heaphood and specific arrangements of grains, one can focus on the relation between the correct applica-

⁸In contrast, the logic of identity does not make appearing to be the same a transitive relation.

⁹It is not claimed that the subject could easily have judged falsely that there is a heap.

tion of the word 'heap' and specific patterns of its use. The problem is now that we cannot recognize the transition between patterns of use in virtue of which the word correctly applies and patterns of use in virtue of which it does not. We may know *that* use determines meaning, but we do not know *how* it does so. In order to identify the determination function exactly for vague words, we should have to identify their correct application exactly, which is just what we cannot do. Whether we experience the use or imagine it, we cannot recognize a stage as a transition. Here too, the epistemic view correctly predicts that we cannot imagine the transition as a transition. That it does so is no objection to the view. When we theorize about vagueness, logic is a better guide than the imagination.

3

The imagination poses another kind of threat to the epistemic view of vagueness. For one may suppose that one can coherently imagine non-epistemic vagueness, by constructing cases in which there is no fact of the matter, in a way that establishes at least that there *could be* non-epistemic vagueness. Given Murphy's Law, the compelling principle that whatever can go wrong will, it follows that there *is* non-epistemic vagueness. The most striking cases involve incomplete stipulations. For example, Kit Fine introduces the predicate 'nice₁' of natural numbers by just the following clauses:¹⁰

- (1) (a) n is nice₁ if $n > 15$.
 (b) n is not nice₁ if $n < 13$.

It is very tempting to say that there is no fact of the matter as to whether 13, 14 or 15 is nice₁. How can the epistemic view be right here? But if it is wrong for incomplete stipulations when they are explicit, as they rarely are, will it not also be wrong for them when they are implicit, as they often are? That is the problem to be confronted here.

One might simply hold that, because the stipulation is incomplete, the term 'nice₁' fails to refer, and sentences containing it fail to say that anything is the case; since the disquotational principles TRUE and FALSE concern only sentences which do say that something is the case, there is no counterexample to the epistemic view. However,

¹⁰See Fine 1975: 266; similar examples are discussed in Williamson 1994: 213-214 and 1995: 189, and used against the epistemic view in Tappenden 1995.

this is an unpromising line for the defence of the epistemic view to take. If nothing is said to the case when the incomplete stipulations are explicit, then presumably nothing is said to be the case when the incomplete stipulations are implicit. But it is plausible to suppose that our language is permeated by what are tantamount to incomplete implicit stipulations; on that supposition, the line implies that our language is permeated by terms which fail to refer and sentences which fail to say that anything is the case. Arguably, this incompleteness pervades even the vocabularies of natural science (Williamson 1994: 169-171). We cannot purify our language by stipulating all the incompleteness away, because the terms in which we should have to make the stipulations would themselves be incomplete. If this view is right, then incoherence is the price of speech. This is not the epistemic view at all, but a kind of nihilism, of intellectual despair.

The defender of the epistemic view should allow that the term 'nice₁' does refer (to the property of being nice₁, of course) and that many sentences involving it do say that something is the case. More cautiously: the problem arises only if we are willing to use the term 'nice₁'; but if we are so willing, then we should also be willing to admit that the sentence '14 is nice₁' says that 14 is nice₁. But if a sentence says that 14 is nice₁, then surely it is true if and only if 14 is nice₁, and false if and only if 14 is not nice₁. Given the law of excluded middle, either 14 is nice₁ or 14 is not nice₁; it follows by elementary logic that the sentence '14 is nice₁' is either true or false. If so, it is obscure what substance there is to the claim that there is no fact of the matter as to whether 14 is nice₁. It is hard to avoid the epistemic view without rejecting classical logic or the disquotational biconditionals for truth and falsity. But if '14 is nice₁' is either true or false, which is it? The case seems so simple that it is very hard to believe that any fact of the matter is deeply hidden from us.

Some more subtle manoeuvres are available to the epistemic view. One could lay down (1)(a)–(b) with the intention that 'nice₁' be understood as a universally quantified variable over all properties satisfying (1)(a)–(b). More precisely, a sentence involving 'nice₁' would be understood as the result of replacing all occurrences of 'nice₁' by a second-level variable 'P' and prefixing the sentence with 'For every P such that n is P whenever $n > 15$ and n is not P whenever $n < 13$ '. On this interpretation, both '14 is nice₁' and '14 is not nice₁' come out as false (not as neither true nor false). There is no mystery; the apparent negation of '14 is nice₁', '14 is not nice₁', is not its real negation, because each sentence is prefixed with an invisible universal quantifier. Classical logic, the disquotational principles TRUE and FALSE and the epistemic view are unthreatened. By

a similar convention, mathematicians sometimes use formulas with free variables to make universally quantified statements. For example, if the variable ' x ' is bound by an implicit universal quantifier over real numbers, then both the inequalities ' $x^2 > x$ ' and ' $x^2 \leq x$ ' are incorrect (false). Our intentions can make such an account right. However, it is implausible to suppose that we often have the requisite intentions, for that would be to attribute a massive complexity of underlying form to our thoughts. In particular, when the incomplete stipulations are merely implicit in our practice, we need not conceptualize them in the way required to think the restrictions on the postulated universal quantifiers (e.g. 'such that n is P whenever $n > 15$ and n is not P whenever $n < 13$ '). Thus the contemplated move is of strictly limited value to the epistemic view. Something more general is needed. In what follows, it will therefore be assumed that '14 is not nice₁' is the real negation of '14 is nice₁'.

The case seems so puzzling because there seems to be perfect symmetry between the hypotheses of truth and falsity; how could anything break the symmetry? But the symmetry is not really perfect. On the epistemic view, the choice in the object-language is between two predicates, 'nice₁' and 'not nice₁'. They are not symmetrical; the second is the negation of the first and not *vice versa*. Either 'nice₁' or 'not nice₁' applies to the unstipulated cases, 13, 14 and 15. Thus either 'nice₁' or 'not nice₁' is *unspecific*, in the sense that it applies to different numbers for quite different reasons: to one by stipulation, to another by some kind of default operative in the unstipulated cases. In principle, 'nice₁' and 'not nice₁' could be simultaneously unspecific; for example, 15 might count as nice₁, 13 and 14 as not nice₁. But given the choice between regarding an atomic predicate as unspecific and so regarding its negation, we generally prefer the latter. For example, we regard 'red' as more specific than 'not red'. When learning a language, we defeasibly assume that an atomic predicate will have this kind of specificity; we do not assume that its negation will. Of course, it is very hard to give a rigorous account of specificity. In Fine's example, 'nice₁' is bound to come out as *less* specific than 'not nice₁' in one sense, however the unstipulated cases are distributed between them, because the former will apply to infinitely many numbers and the latter to only finitely many. But that is consistent with this (vague) default principle:

SPEC All other things being equal, the application conditions of an atomic predicate are as specific as possible.

SPEC is to be understood as a constitutive principle for application conditions, not just as a rule of thumb that we use in language learn-

ing. Of course, all other things hardly ever are equal: but the persuasive force of examples such as Fine's depends on the impression that in this case they are equal, for that *is* the symmetry problem. When SPEC is applied to 'nice₁', it implies that 'nice₁' does not apply to the unstipulated cases. On the epistemic view, it follows that 13, 14 and 15 are not nice₁, so 'not nice₁' does apply to them. Some numbers fail to be nice₁ because they were stipulated not to be nice₁; others fail to be nice₁ because they were left unstipulated. SPEC breaks the symmetry.

A term might be introduced by an incomplete stipulation with the intention that principles like SPEC should be irrelevant to its application. If the intention is fulfilled, then sentences involving the term will presumably have the kind of universally quantified sense discussed earlier.

The epistemic view of vagueness does not entail SPEC. What matters is that the two are consistent, so that their conjunction can give a defensible account of examples like 'nice₁'. If the account is right, and we can know that it is, then we can know that 'nice₁' does not apply to the unstipulated cases. Thus, on the epistemic view, 'nice₁' is not a good example of a vague predicate. That is plausible anyway; the precise and explicit stipulation (1)(a)-(b) is quite unlike anything which governs a vague predicate such as 'red'. The unsurveyable complexity of the pattern of our use, central to the account of unknowability in §II, is absent by hypothesis. What remains is an epistemic difference between the stipulated and the unstipulated cases in the conditions for understanding 'nice₁': to do so, one must know that if $n < 13$ then 'nice₁' does not apply to n (or perhaps one must know that *it is stipulated that* if $n < 13$ then 'nice₁' does not apply to n); one need not know that if $13 \leq n \leq 15$ then 'nice₁' does not apply to n . One must know the stipulation; one need not know its effects.

Note that 1(a)-(b) has been reformulated in metalinguistic terms, with 'nice₁' in quotation marks, so that the connection between stipulation and understanding can be properly stated. Someone could know that n is nice₁ whenever $n > 15$ and that n is not nice₁ whenever $n < 13$ by understanding a synonym of 'nice₁' without knowing anything at all about the term 'nice₁' itself. It is independently plausible that a stipulation should be formulated in terms whose meaningfulness does not require the stipulation already to have been made. 1(a)-(b) will henceforth be taken in this metalinguistic way.

Does the envisaged account allow a sense in which it has been *left open* whether 'nice₁' applies to the unstipulated cases? Suppose that one subsequently stipulates that 'nice₁' will apply to the unstipulated cases. On the envisaged account, the application condition for 'nice₁'

ceases to be that $n > 15$; it is now that $n \leq 13$. Thus ‘nice₁’ has changed in meaning. But everyone thinks that; it has certainly changed in meaning if ‘14 is nice₁’ has changed from being half true or neither true nor false to being true. The question is whether the new meaning can somehow be regarded as a *completion* of the old one. On the epistemic view, it can in this way: the stipulation which one must know in order to understand ‘nice₁’ in its new sense completes the stipulation which one must know in order to understand ‘nice₁’ in its old sense.

Although the new stipulation reverses the truth-value of some sentences, that does not make it *mistaken*, for the stipulator need not have been trying to preserve the truth-value of all sentences; the intention may have been merely to preserve the original stipulation. The point of sharpening is to introduce clear application conditions which preserve the clear cases of the old application conditions, whether or not clarity is an epistemic matter. When we are free to sharpen, the old truth-values in the unclear cases are not worth preserving.

What if we stipulate instead that ‘nice₁’ does *not* apply to the unstipulated cases 13, 14 and 15?¹¹ The extension of ‘nice₁’ has not changed. But if meaning determines understanding, and understanding requires knowledge of the relevant stipulations, then the meaning of ‘nice₁’ has changed. For after the stipulation that ‘nice₁’ does not apply to 13, 14 and 15, understanding ‘nice₁’ requires knowing that ‘nice₁’ does not apply to them (or perhaps knowing that it is stipulated that ‘nice₁’ does not apply to them). Before the stipulation, that knowledge was not required for understanding. Thus the condition for understanding ‘nice₁’ has changed. If meaning determines understanding, in the sense that any difference in the conditions for understanding expressions implies a difference in their meanings, then it follows that the meaning of ‘nice₁’ has changed.

SPEC may appear not to help with some incomplete stipulations. Consider, for example, this disjunctive stipulation for a predicate ‘edd’ of natural numbers:

- (2) Either ‘edd’ applies to all and only even numbers, or ‘edd’ applies to all and only odd numbers.

Thus ‘edd’ has two possible application conditions, of equal specificity: that n is even, and that n is odd. SPEC does not break the symmetry between those two conditions. Of course, stipulations

¹¹Michael Tye raised this question.

such as (2) are not made in practice, but that point by itself does not constitute a very satisfying defence of the epistemic view. Fortunately, SPEC can be applied in a different way. Given just (2), the grounds for 'edd' to apply to any one natural number are exactly as strong as the grounds for it to apply to any other. Thus the salient hypotheses are that 'edd' applies to all natural numbers and that it applies to none. SPEC does decide between these two hypotheses, in favour of the latter: nothing is edd. Of course, both hypotheses are inconsistent with (2), which entails that 'edd' applies to some but not all natural numbers. That would be a problem if stipulating that P always made it true that P: but stipulations have no such magic effect. The most obvious case in point is an inconsistent stipulation. Similarly, to stipulate that a new term is meaningful, without stipulating what it means, is not even to make that term meaningful. To stipulate (2) is not thereby to make (2) true. If understanding 'edd' with the sense conferred by (2) entails knowing the truth of (2), then one cannot understand 'edd' with that sense. Perhaps, however, one can understand 'edd' just by knowing that the stipulation is (2), without knowing the truth of (2) (in this way one could understand terms governed by inconsistent stipulations). Either way, the crucial gap is between stipulating that P and making it true that P.

The gap between stipulating and making true emerges in a different way for 'nice₁'. The content of the stipulation for 'nice₁' does not entail that 14 is not nice₁. Nevertheless, in the absence of countervailing factors, the stipulation made it true that 14 is not nice₁.

When a stipulation that P fail to makes it true that P, it is not as though nothing had been said. The stipulation may remain 'on the books', to be reactivated by later stipulations. For example, once (2) has been stipulated, it can later be stipulated that 'edd' applies to 7, which breaks the symmetry between the disjuncts of (2). The two stipulations together might suffice to make 'edd' apply to all and only odd numbers. The later stipulation would not have had this effect by itself.

The effect of a stipulation is not determined by its content, but depends also on the intentions and authority with which it is made, and the understanding and deference with which it is received. When the different cases are properly separated, no counterexample to the epistemic view can be found.¹²

¹²I thank audiences at the University of St Andrews and the SOFIA conference, and in particular Mario Gómez-Torrente, Terry Horgan and Michael Tye, for helpful comments, and the University of Edinburgh and the Carnegie Trust for the Universities of Scotland for financial support.

BIBLIOGRAPHY

- Fine, Kit. 1975. 'Vagueness, truth and logic'. *Synthese* 30: 265-300.
- Fodor, Jerry, and LePore, Ernest. 1996 'The red herring and the pet fish: why concepts still can't be prototypes'. *Cognition* 58: 253-270.
- Kamp, Hans, and Partee, Barbara. 1995. 'Prototype theory and compositionality'. *Cognition* 57: 129-191.
- Lewis, David. 1986. *On the Plurality of Worlds*. Oxford: Blackwell.
- McGee, Vann, and McLaughlin, Brian. 1995. 'Distinctions without a difference'. *The Southern Journal of Philosophy* 33, Supplement: 203-251.
- Osherson, Daniel, and Smith, Edward. Forthcoming. 'On typicality and vagueness'. *Cognition*.
- Priest, Graham. 1987. *In Contradiction*. Dordrecht: Nijhoff.
- Tappenden, Jamie. 1995. 'Some remarks on vagueness and a dynamic conception of language'. *The Southern Journal of Philosophy* 33, Supplement: 193-201.
- Williamson, Timothy. 1994. *Vagueness*. London: Routledge.
- . 1995. 'Definiteness and knowability'. *The Southern Journal of Philosophy* 33, Supplement: 171-191.