

5 Intuitionism and the Sorites Paradox

Crispin Wright

5.1 Introduction – the Basic Analogy

Mathematical platonism may be characterised as the conviction that in pure mathematics we explore an objective, abstract realm that confers determinate truth-values on the statements of mathematical theory irrespective of human (finite) capacities of proof or refutation. This conviction crystallises in the belief that classical logic, based on the semantic principle of bivalence, is the appropriate logical medium for pure mathematical inference even when, as of course obtains in all areas of significant mathematical interest from number-theory upwards, we have no guarantee of the decidability by proof of every problem. In this respect – the conviction that the truth-values, *true* and *false*, are distributed exhaustively and exclusively across a targeted range of statements irrespective of our cognitive limitations – an epistemicist conception of vagueness¹ bears an analogy to the Platonist philosophy of mathematics. Let us characterise a vague predicate as basic just if it is semantically unstructured and is characteristically applied and denied non-inferentially, on the basis of (casual) observation. The usual suspects in the Sorites literature – *bald*, *yellow*, *tall*, *heap* and so on – are all of this character and will provide our implicit focus in what follows. Epistemicism postulates a realm of distinctions drawn by such basic vague concepts that underwrite absolutely sharp ‘cut-offs’ in suitable soritical series,² irrespective of our capacity to locate them. For the epistemicist, the principle of bivalence remains good for vague languages – or if it does not, it is not vagueness that compromises it – and classical logic remains the appropriate medium of inference among vague statements. An indiscernible difference between two colour patches in a soritical series for *yellow* may thus mark an abrupt transition from yellow to orange; the impression of the *indeterminacy* of that distinction is merely a reflection of our misunderstanding of our ignorance of where the cut-off falls. For the epistemicist, the Sorites Paradox is accordingly easily resolved. It is scotched by the simple reflection that its major premise will always be subject to counterexample in any particular soritical series. If the initial element is yellow and the

¹ As supported by, among others, Cargile (1969), Horwich (1998b), Sorensen (1988) and, in its most thoroughgoing development, Williamson (1994). See also Chapter 1 in this volume.

² We understand a ‘suitable’ Sorites series for a predicate *F* to be *monotonic*, that is, one in which any *F*-relevant changes involved in the move from one element to its immediate successor are never such that the latter has a stronger claim to be *F* than the former.

final element is orange, then there must be an adjacent pair of elements, one of which is yellow while the next is orange. It is just that we, in our ignorance, are in no position to identify the critical pair.

On one understanding of it, the *ur*-thought of intuitionism as a philosophy of mathematics is a rejection of the idea of a potentially proof-transcendent mathematical reality as a *superstition*: something that there is, simply, no good reason to believe in. For the intuitionist, the mathematical facts are justifiably regarded as determinate only insofar as they are determinable by proof, and the relevant notion of proof needs accordingly to be disciplined in such a way as to avoid any implicit reliance on the platonist metaphysics. So, in any area of mathematics where we lack any guarantee of decidability, the logic deployed in proof construction cannot rely on the principle of bivalence and hence – according to the intuitionist – cannot justifiably be classical. In particular the validity of the law of excluded middle, which intuitionism understands as depending on the soundness of bivalence, can no longer be taken for granted. There is evident scope for a similar reaction to the epistemicist conception of vagueness. The latter is a commitment to a transcendent semantics for vague expressions, which construes them as somehow glomming onto semantic values – properties in the case of vague predicates – that are possessed of absolutely sharp extensions, potentially beyond our ken. The conception that vague expressions work like that may likewise impress as the merest superstition. Perhaps a little more kindly, it may impress as merely *ad hoc*, for there is not the slightest reason that speaks in favour of it except its convenience in the context of addressing the Sorites.³ Say that an object, *o*, is *F*-surveyable just in case *o* is available and open to as careful an inspection as is necessary to justify the application of *F* to it whenever it can be justified. For the epistemicist, reality is such that the application of any meaningful basic vague predicate, *F*, to an *F*-surveyable object must result in a statement of determinate truth value, true or false. For an intuitionistic conception of vagueness – one conceived on the model of mathematical intuitionism – a satisfactory semantics and logic for basic vague predicates must eschew commitment to any such claim.

The avoidance of such a commitment is of course common ground with any instance of the long tradition of theories about vagueness that construe borderline cases as examples of *semantic indeterminacy*: as cases where the rules of the language leave us in the lurch, so to speak, by issuing no instruction for any particular verdict. Here, though, the intuitionist credits the epistemicist with a crucial insight: that vagueness is indeed a *cognitive* rather than a semantic phenomenon; that our inability to apply the concepts on either side of a vague distinction with consistent mutual precision is not a *consequence* of some kind of indeterminacy or incompleteness in the semantics of vague expressions, but is constitutive of the phenomenon.

³ I have sometimes encountered in discussion the impression that the motivational shortfall here is addressed by Williamson's argument that knowledge everywhere requires a margin of error. Not so. What that argument establishes, if anything, is that if there is a sharp cut-off in a Sorites series, we will not be able to locate it. The argument provides no reason to suppose that the antecedent of that conditional is true. I'll say a little more about the Williamsonian argument below. See also Chapter 1 in this volume, Sections 1.2 and 1.3.

Consider this example. Suppose we are to review a line of 100 soldiers arranged in order of decreasing height and to judge of each whether he/she is at least 5 ft 10 in tall – but to judge by eye rather than by using any means of exact measurement. Let the soldiers' heights range from 6 ft 6 in to 5 ft 6 in. This provides a toy model of a Sorites series as conceived by the epistemicist. For while there is indeed a sharp cut-off – there must be a first soldier in the line who is less than 5 ft 10 in tall – our judgements about the individual cases will expectably divide between an initial range of confident positive verdicts and a later range of confident negative verdicts between which there will be a region of uncertainty, where we return hesitant, sometimes mutually conflicting verdicts and sometimes struggle to return a verdict at all. Here, of course, we have a conception of canonical grounds for determining whether a hard case has the property expressed by the predicate at issue, so that is a point of contrast with our situation when we face a Sorites series for a vague predicate as conceived by epistemicism. But putting that disanalogy to one side, it remains that in the soldiers scenario, our patterns of judgement will have exactly the physiognomy that the epistemicist regards as the hallmarks of vagueness. Hence, in her view, there is nothing in our practice with vague concepts that distinguishes it from judgements concerning sharply bounded properties about whose specific nature we are ignorant.

The intuitionist agrees with epistemicism that such a physiognomy of practical judgement is characteristic of vagueness. But intuitionism drops both the assumption that the judgements concerned are answerable to the extension of a sharply bounded property *and* the notion that a different kind of explanation of these characteristic judgemental patterns is called for, in terms, roughly, of shortcomings in – our lack of guidance by – the semantic rules that fix the meanings of the expressions concerned. For the intuitionist, the vagueness of a predicate *consists in* these distinctive patterns in our use of it. They are the whole story. The intuitionist conception of vagueness is thus a *deflationary* conception: it holds that there is no more to the phenomenon than meets the eye, so to speak – that it is unnecessary, is indeed a mistake, to look to some underlying feature of the semantics of vague expressions to explain our characteristic patterns of judgement in the borderline area. (That is not to say that one should not look for an explanation of a different kind.) It is the view of the intuitionist that both epistemicism and indeterminism commit versions of this mistake.

The justification for this charge when the canvassed alternative is semantic indeterminism rests on a complex variety of considerations whose details, for reasons of space, I cannot rehearse here.⁴ However there is one such consideration – one relevant aspect of our practice with vague concepts – which is particularly important for the grounding of the most distinctive aspect of the intuitionist approach. Semantic indeterminism interprets borderline cases of a distinction as cases where there is no mandate to apply either of the expressions concerned – where the rules for their use prescribe no verdict. That suggestion does a poor job of predicting one salient aspect of our judgemental practices with vague concepts,

⁴ For some elaboration see Wright (2003a, 2007, 2010); also Williamson (1994).

namely our uncritical attitude to polar – positive or negative – judgements concerning items in their borderline regions. Provided a verdict is suitably qualified and evinces an awareness that the case is a marginal one, it is not treated as a mark of incompetence, or mistake, to have a view, positive or negative, about any single borderline case.⁵ Suppose *X* struggles to have an opinion whether some shade from the mid-region of the yellow-orange Sorites is yellow enough to count as yellow but *Y* is of the opinion that it is – just about – yellow. Our sense is that such divergences are just what is to be expected, and that each reaction can be as good as the other. *X* need not be regarded as coming short; *Y* need not be regarded as overreaching. Each reaction is quite consistent with full mastery of *yellow* and due attention to the hue concerned.

According to semantic indeterminism, this *laissez faire* attitude should be regarded as cavalier, for *X* and *Y* cannot both be operating as the relevant semantic rules require; the rules cannot both be silent on the relevant hue and mandate *Y*'s qualified verdict of 'yellow'.⁶ Yet our ordinary practice reflects no sense of that. We are characteristically open to – as I have elsewhere expressed it (Wright 2003a and 2007), *liberal* about – polar verdicts about borderline cases. To be sure, the indeterminist might be tempted to interpret this liberality as reflecting a sense of respect for our ignorance about in just which cases the rules do in fact fall silent – which are the true borderline cases. But since, if so, there is no evident means of remedying that ignorance, that again would be a step in the direction of objectionably transcendentalising the semantics of vagueness. For the intuitionist, in contrast, there really need be no sense in which one who returns a (qualified) polar verdict about a borderline cases does worse than one who fails to reach a verdict.

The point may seem slight, but it is crucial. For respecting this aspect of our practice as in good standing requires that, in contrast to the view of semantic indeterminism, we should not regard borderline cases as presenting *truth-value gaps*.⁷ If borderline cases are truth-value gaps, then someone who returns a polar verdict about such a case actually makes a mistake. And that is just what, according to liberalism, we have no right to think. It follows that we have no right to regard borderline cases as *counterexamples* to the principle of bivalence, and hence that vagueness, as now understood, provides no motive to *reject* bivalence. Since, by rejecting epistemicism, and recognising that we cannot in general settle questions in the borderline region either, we have also undercut all motive to *endorse* the

⁵ For the purposes of this claim, we may take a borderline case to be any that tends to elicit the judgemental physiognomy characterised earlier among a significant number of competent judges.

⁶ To be sure, there is another possibility: we might try to think of the rules as, in the borderline area, issuing *permissions*. Then both a tentative verdict and a failure, or unwillingness, to reach a verdict, may be viewed as rule-compliant. But it is very doubtful that any satisfactory proposal lies in this direction. Presumably among the clear cases the rules must *mandate* specific verdicts rather than merely permit them. So we need to ask about the character of the transition from cases where a positive verdict about *F* is mandated to cases where it is merely permitted. If this is a sharp boundary, then since there is again no possibility of knowing where it falls, the proposal will have 'transcendentalised' the semantics of *F* in a manner different from but no less inherently objectionable than epistemicism. But if the transition is accomplished by a spread of further borderline cases – cases that are borderline for the distinction between 'mandatorily judged as *F*' and 'permissibly judged as *F*' – then the question arises how, in point of mandate or permission, cases in this category are to be described. For argument that contradiction ensues, see Wright (2003a).

⁷ Or indeed as having any kind of 'Third Possibility' status inconsistent with each of truth and falsity *simpliciter*. For further discussion, see Wright (2001).

principle, the resulting position is exactly analogous to the attitude of the mathematical intuitionist to bivalence in mathematics: that it is a principle towards which we should take an agnostic stance.

With these preliminaries in place, let us turn to review how an intuitionistic treatment of the Sorites may be developed in more detail.

5.2 The Tolerance and ‘No Sharp Boundaries’ Paradoxes

The classic deductive⁸ Sorites paradoxes vary in two respects: first in the formal character of the major premise involved, and second – where the major premise is shared – in the manner in which that premise is made to seem plausible. And of course different forms of major premise will call for correspondingly different deductive sub-routines in the derivation of the paradox. Perhaps the most familiar form of the deductive Sorites is what we may call the Tolerance Paradox.⁹ As normally formulated, it uses a universally quantified conditional major premise:

TP: $(\forall x)(Fx \rightarrow Fx')$

and proceeds on the assumption of one polar premise, $F1$, and $n - 1$ successive steps of universal instantiation and *modus ponens* to contradict the other polar premise, $\sim Fn$. As for motivation, the key thought is, as the title I have given to the paradox suggests, that, such is its meaning, the application of F , and/or the justification for applying it, *tolerates* whatever small changes may be involved in the transition from one element of the series to the next: for instance, that if a colour patch is (justifiably described as) red, a pairwise indiscriminable (or even just barely noticeable) change in shade must leave it (justifiably described as) red; that if a man is bald, the addition of a single hair won’t relevantly change matters, and so on. For the examples with which we are concerned, claims of this ilk can seem thoroughly intuitive; and they can be supported by a variety of serious-seeming theoretical considerations.¹⁰ In some cases, indeed, the claim of tolerance may seem absolutely unassailable: how could ‘looks red’, for example, fail to apply to both, if to either, of any pair of items that look exactly the same? Unfortunate, then, that ‘looks exactly the same’ is not a transitive relation.

The Tolerance Paradox, however, impressive as it may be in particular cases, is not, or at least not obviously, a paradox of vagueness *per se*. Vagueness is not, or at least not obviously, the same thing as tolerance. Precision must imply non-tolerance, of course, but the converse is intuitively less clear. Ought there not somehow to be some distance between a predicate’s possession of borderline cases and its being tolerant of some degree of marginal change? While the claim may indeed seem intuitive, it requires – in the presence

⁸ As distinguished from the so-called *Forced March Sorites*. For reasons of space, I must forgo discussion of that here. See ‘The Forced-March Paradox’, in the introduction to this volume.

⁹ See also ‘The Sorites Paradox’, in the introduction to this volume.

¹⁰ For elaboration of some such, see Wright (1975).

of paradox – argument to suggest that *yellow, heap, bald*, etc., are tolerant. But no argument is required to suggest that they are vague. That these predicates are vague is a *datum*.

The No Sharp Boundaries Paradox, by contrast, impresses as a paradox of vagueness par excellence. It works with a negative existential major premise,

NSB: $\sim (\exists x)(Fx \ \& \ \sim Fx')$

that may very plausibly seem simply to give expression to what it is for F to be vague in the series of objects in question. For vagueness, surely, is just the complement of precision, and the sentence of which that negative existential is the negation, viz. what I have elsewhere (Wright 2007) called the *unpalatable existential*

UE: $(\exists x)(Fx \ \& \ \sim Fx')$,

surely just *states* that F is precise in the series in question: that there is a sharp boundary between the F s and the non- F s, and so no borderline cases. If, then, F is in fact vague, the negative existential seems imposed just by that fact, indeed to be a statement of exactly that fact. And now contradiction follows by iteration of a different but no less basic and cogent-seeming deductive sub-routine, involving conjunction introduction, existential generalisation, and *reductio ad absurdum* as a negation introduction rule.¹¹

With both paradoxes, there is the option of letting the reasoning stand as a *reductio* of the major premise. If we take that option with the Tolerance Paradox, we treat it as a schematic proof that none of the usual suspects *is* genuinely tolerant of the marginal differences characteristic of the transitions in a soritical series for it. Tolerance, in that case, is simply an illusion. And that is a conclusion we might very well essay to live with, provided we can provide a satisfactory explanation of why and how the illusion tends to take us in, and of what is wrong with the ‘serious-seeming theoretical considerations’ apparently enforcing tolerance that I have already alluded to.

But not so fast: even if those obligations can be discharged, the proposed response, in the presence of classical logic, is not yet stable. For (allowing its ingredient conditional to be material) the negation of TP, now regarded as established by the paradoxical reasoning, is a classical equivalent of the unpalatable existential. So if our logic is classical, non-tolerance does after all collapse into precision, and to the extent that one feels there should, as remarked above, be daylight between them, that should impress as a black mark against classical logic in this context. Moreover that impression is only reinforced when one considers the option of letting the No Sharp Boundaries Paradox stand as a refutation of NSB. For then all that stands between that result and affirmation of the unpalatable existential is a double negation elimination step. And now, once constrained by classical logic to allow the inference to UE, we seem to be on the verge of admitting that *vagueness itself* is an illusion. That, surely, isn’t anything we can live with.

¹¹ That is, the intuitionistically valid half of classical *reductio*, where the latter also allows *reductio ad absurdum* inferences that serve to eliminate negations.

Intuitionism, by contrast, aims at winning through to a position where we can accept each of the Tolerance and No Sharp Boundaries Paradox as a *reductio* of its major premise but refuse in a principled way the inference onwards to the unpalatable existential. We also aim to retain the ordinary conception of an existential statement as requiring a determinate witness for its truth, and thus to avoid any form of the implausible semantic story that construes the statement ‘There is in this series a last *F* element followed immediately by a first non-*F* one’ as neutral on the question of the existence of a sharp cut-off as intuitively understood. Our path will be to explore, in the light of the general, deflationary conception of the nature of vagueness outlined earlier, what motivation it may be possible to give for broadly intuitionist restrictions on the logic of inferences among vagueness-involving statements. In this we follow a suggestion first briefly floated at the end of Hilary Putnam (1983).¹² If, in particular, we can justify a rejection of double negation elimination for molecular vague statements in general, then it may be possible comfortably to acknowledge that both the Tolerance and the No Sharp Boundaries Paradox do indeed disprove their respective major premises without any consequent commitment to the unpalatable existential, nor consequent obligation either (with the epistemicist) to believe it or (with the supervaluationist) to somehow reinterpret it in such a way that it doesn’t mean what it seemingly says.

5.3 Constraints on an Intuitionistic Solution

I propose that we set the following three constraints on the project. First (*Constraint 1*) and most obvious, we need to *motivate* the required restrictions on classical logic in general and, in particular, to explain how a valid *reductio* of TP, or NSB, can fail to justify the unpalatable existential.

Second (*Constraint 2*), as with all attempts to solve rather than merely block a paradox, we must offer a convincing explanation of why the premises that spawn aporia impress us as plausible in the first place, of what mistaken assumptions we have implicitly fallen into that give them their spurious credibility. So in the present instances, we must contrive to explain away the continuing powerful temptation to regard the major premises for the Tolerance and No Sharp Boundaries Paradox as true. I have said much elsewhere to attempt to defuse the attractions of tolerance premises.¹³ Here we will focus on the challenge to explain why NSB is *not* a satisfactory characterisation of *F*’s vagueness in the series in question. (We have already implicitly shown our hand on this.)

Finally, (*Constraint 3*) I think it reasonable to require, although I grant it is not wholly clear in advance exactly what the requirement comes to, that Constraints 1 and 2 should, so far as we can manage it, be satisfied in a way that draws on an overarching account of

¹² Early discussions of Putnam’s proposal, besides my own work from (2001) onwards, include Read and Wright (1985), Putnam (1985), Schwartz (1987), Rea (1989), Schwartz and Throop (1991), Putnam (1991), Mott (1994), Williamson (1996b) and Chambers (1998).

¹³ Such an attempt must perforce be somewhat ramified, in order to match the diverse sources of such attraction. My own diagnostic forays run from Wright (1975) through Wright (1987) to Wright (2007).

what the relevant kind of vagueness consists in (i.e. of the nature of the relevant kind of borderline cases.) We are proposing restrictions on what, from a classical point of view, are entrenched, tried and tested patterns of inference. If such restrictions are justified, it may be, to be sure, that that justification is global, applying within discourses of every kind. That is the character, for example, of the meta-semantic considerations about acquisition and manifestation of understanding originally offered by Michael Dummett half a century ago in support of a global repudiation of the principle of bivalence except in areas where decidability is guaranteed. Whatever one's estimate of such arguments, what Constraint 3 is seeking is a justification for relevant restrictions on classical logic that is specifically driven by aspects of the nature of vague discourse. In the present context, that will require putting to work the deflationary conception of vagueness sketched in Section 1.

5.4 Addressing Constraint 1: The Basic Revisionary Argument

In the mathematical case, as remarked, the intuitionistic attitude flows from a rejection of the principle of bivalence, based on a repudiation of platonist metaphysics and insistence that truth in pure mathematics can only consist in the availability of proof.¹⁴ In the case of vague statements, many would be pre-theoretically willing to grant that bivalence is generally unacceptable anyway. Certainly, the metaphysics of meaning implicit in epistemicism has none of the intuitive appeal of, say, arithmetical platonism. But even if it is granted that bivalence is *principium non gratum* where vagueness is concerned, repudiating the principle is one thing and motivating revision of classical logic a further thing. Classical logic need not necessarily fail if bivalence is dropped. How should the intuitionist argue specifically that the *logic* of vague discourse should not be classical?

What I once called the 'basic revisionary argument' is designed to accomplish that result. It runs for any range of statements that are not guaranteed to be decidable but are subject to a pair of principles of *evidential constraint* (EC). That is, for each such statement *P*, each of these conditionals is to hold:

EC: $P \rightarrow$ it is feasible to know *P*

Not $P \rightarrow$ it is feasible to know Not-*P*

Now, it is plausible – but with caveats, to be considered in a moment – that each of the usual suspects (*yellow, bald, tall, heap, etc.*) generates atomic predications that exhibit this form of evidential constraint; that is, intuitively, if something is, in the sense characterised earlier, surveyable for *yellow* (that is, it is available for inspection in decent conditions, etc.), and it *is* yellow, then we'll be able to tell that it is; and if it *isn't* yellow, we'll be able to tell that. Intuitively, what colour something is *cannot hide* if and when conditions present themselves in which it is possible to have a proper look at it.¹⁵ And analogously for baldness, tallness and 'heaphood'. The basic revisionary argument is then the observation

¹⁴ This argument is central in Wright (1992, 2001, 2007).

¹⁵ The claim that EC holds good for these cases is thus not subject to 'killer yellow' issues.

that, if the law of excluded middle is retained for all such predications, P , then reasoning by cases across the EC-conditionals will disclose a commitment to the disjunction:

D: It is feasible to know $P \vee$ It is feasible to know not- P

In effect, the thesis that P is decidable. But of that, if the relevant predicate is associated with borderline cases, we have no guarantee. Accordingly we have no guarantee of the validity of the law of excluded middle in application to such statements and therefore have no business reckoning it among the logical laws.

Simply expressed, the thrust of the argument is that a range of statements may be such as both to *lack* any general guarantee of decidability in an arbitrary instance and to *have* a guarantee that if any of them is true, it will be recognisable true and, if false, recognisably false. Imposition of the law of excluded middle onto such statements will then enforce the conclusion that each of them is decidable true or false – contrary to hypothesis. It will amount to the pretence of a guarantee that we do not actually have.

Arguably a very large class of statements are in this position, including not merely vague predications but, for instance, evaluations of a wide variety of kinds, including expressions of personal taste, humour and perhaps (some aspects of) morality. And of course the argument will run for any region of discourse where we reject the idea that truth can outrun all possibility of recognition but have to acknowledge that we lack the means to decide an arbitrary question – exactly the combination credited by the intuitionists for number-theory and analysis.

Suppose then that we disdain the law of excluded middle on this (or some or other) basis. The soritical series we are considering involve a *monotonic* direction of change: that is, any F element is preceded only by F elements, and any non- F element is succeeded only by non- F elements. The reader will observe, accordingly, that once the law of excluded middle is rejected, the sought-for distinction between the unpalatable existential and its double negation is enforced. For the latter is surely established by the inconsistency of NSB with the truths expressed by the polar assumptions. But, given monotonicity, the unpalatable existential is equivalent to the law of excluded middle over the range of atomic predications of F on the series of elements in question.

5.5 One Objection to the Basic Revisionary Argument

So far, so good. But now for the caveats. The EC-conditionals are challengeable on a number of serious-looking counts. First, they are in direct tension with the upshot of Timothy Williamson's recently influential 'anti-luminosity' argument.¹⁶ Familiarly, Williamson makes a case that if knowledge generally is to be subject to a certain form of (putatively) plausible safety constraint, then it must be controlled by a margin of error: in particular, if a subject knows that F holds of an object a , it cannot be that F fails to hold of any object

¹⁶ Williamson (2000, chapter 4).

that she could not easily distinguish, using the same methods, from a . The effect is thus that, for elements, x , in a soritical series for F , the following conditional is good:

(It is feasible to know that Fx) $\rightarrow Fx'$,

which, paired with the first of the EC-conditionals, immediately provides the means to show that F applies throughout the soritical series.

Here is not the place for a detailed engagement with Williamson's thesis. But there are a couple of fairly immediate misgivings about it that deserve notice. One is whether the notion of safety that it utilises is indeed a well-motivated constraint on knowledge everywhere, whatever the subject matter and methods employed. Williamson's intuitive thought, if I may venture a precis,¹⁷ is that if a subject comes to the judgement that Fa , and a' is pairwise indistinguishable from a , then the subject must be significantly likely also to judge that Fa' – and now, if the latter is false, she is therefore very likely to make a false judgement using the very methods she used in judging Fa . So those methods are not generally reliable, in which case the judgement that Fa , based upon them, ought not to count as knowledgeable in the first place. Yet if that is the intuitive thought, one salient question is why we should require that, in order to count as a reliable means for settling a question about one item, a method must also be reliable about *others* that, however similar, differ from it. Why could not a machine – a speedometer, for example – that issues a varying digital signal in response to a varying stimulus have an absolutely sharp threshold of reliability, so that its responsive signals are reliable up to and including some specific value, k , in its inputs but then go haywire for inputs of any greater value. In that case, its signals may be regarded as 'knowledgeable' for any input value i , less than but as close to k as you like. If it is not a priori ruled out that our judgements, for some particular pairings of subject matter and methods, are like that, then it is not a priori guaranteed that Williamsonian safety is everywhere a necessary condition of knowledge.

One can envisage the likely rejoinder that as a matter of anthropological fact we are not in any area of our cognitive activities comparable to such a machine. Still, even if that is so, it seems incredible that such a contingency could somehow entail that there are yellows, and instances of baldness, etc., that lie beyond our powers of recognition even in the best of circumstances.

But now grant that the general requirement of safety proposed – again, the proposal that in order to know that P in circumstances C my methods must be such that they could not easily lead me astray in circumstances sufficiently similar to C – grant that this is well motivated everywhere. A second misgiving is that in the way that Williamson puts the proposal to work, no account is taken of the possibility of *response-dependence*: the idea that some kinds of judgement – and here the critic is likely to be thinking of exactly the kinds of judgement, about sensations and other aspects of one's occurrent mental state, that Williamson means to target in directing his argument against the traditional idea of our

¹⁷ Compare Williamson (2000, p. 97).

'cognitive home' – are not purely discriminatory of matters constituted independently but are such, rather, that the subject's own judgemental dispositions are somehow themselves implicated in the facts being judged. For any area of judgement where this idea has traction, the supposition that in perfectly good conditions of judgement we might easily respond to what is in fact a non-*F* case in the way we do to an *F* case that is very similar to it is in jeopardy of incoherence. Simply, if *F*-ness and non-*F*-ness are response-dependent matters, then it cannot legitimately be assumed that, purely on the basis of their similarity in a particular case, we will be at risk of responding to a non-*F* case in the way we do to an *F*-case.

To be sure, the heyday of the recent discussion of response-dependence has passed, and rigorous but still dialectically useful formulations of it proved hard to come by when the debates were at their height. Still, many may feel that there is an elusive truth in it, with qualities instantiated in one's phenomenal mental life and Lockean secondary qualities of external objects generally providing two examples of domains to which philosophical justice can be done only by keeping a place for the idea of response-dependence on our philosophical agenda.¹⁸

I do not think, accordingly, that Williamson's argument, in our present state of understanding, comes anywhere near to establishing that the basic revisionary argument is hobbled by its reliance on the EC-conditionals. Rather the argument sets up yet another philosophical paradox: prima facie plausible thoughts about knowledge in general and a putative requirement of the safety/reliability of methods whereby beliefs are formed, prove to conflict with prima facie plausible thoughts about the luminosity of a range of concepts for which, we would probably otherwise be inclined to think, the EC-conditionals look good. Something has to give. But here cannot be the place to further investigate what.

5.6 Two Further Objections

There are, however, two less theoretically loaded reservations about the role of the EC-conditionals in the basic revisionary argument that should be tabled when what is envisaged is its application specifically to vague expressions. First, no connection has actually been explained to link the EC-conditionals with vagueness as such. All that has been offered is the suggestion that the conditionals are plausible for some examples of vague predicates – for the usual suspects. A general theoretical connection is wanted before there can be any firm prospect of a solution by this route to the Sorites Paradox in general. One senses that a development may be possible of a general connection between vague judgement and response-dependent judgement, grounded in the thought that the status of something as a borderline case is a response-dependent matter. That suggestion, though, once again in the present state of our understanding, is merely speculative.

¹⁸ Concerns of this character, although he does not mention the notion, response- (or judgement-) dependence by name nor relate his discussion to the literature about it, are nicely elaborated in Berker (2008).

Second, and perhaps more threatening to this particular strategy for underwriting an intuitionistic treatment of the Sorites, is the conflict between the EC-conditionals and a principle I have elsewhere called *Verdict Exclusion* (VE):

VE: Knowledge is not feasible about borderline cases.

EC and VE are pairwise inconsistent (since, as the reader will speedily see, they combine to enforce contradictory descriptions of borderline cases.) So someone who accepts EC must deny VE. But VE may well impress – indeed has impressed a number of expert commentators¹⁹ – as a datum. In any case, the principle may seem to have powerful intuitive support from the very deflationary conception of vagueness which, I have proposed, should be seen as the mainspring of an intuitionistic treatment. On that conception, borderline cases are constitutively cases whereby subjects characteristically fall into weak, inconstant and mutually conflicting opinions. Any opinion a subject holds about such a case is one that she might very easily, using just the same belief-forming methods, not have held. Surely on any reasonable interpretation of a safety, or reliability, constraint on knowledge, that must count as inconsistent with such an opinion's being knowledgeable.

Elsewhere, I have suggested that an endorsement of VE proves, on closer inspection, to be in tension with liberalism.²⁰ Let me here make a different point. Once it is given that something is a borderline case, I think the line of argument just outlined for VE is likely to prove compelling. But the crucial consideration is that, of any particular element in a Sorites series, it is *not* a given – except as a contingent point about the sociology of a particular group of judges – that it is a borderline case. Being a borderline case is judge-relative: *x* may be such as to induce the characteristic judgemental difficulties and variability in some but not other competent judges. Let the proposition that *x* is yellow elicit those characteristic responses in some of us but suppose that Steady Freddy consistently judges *x* yellow (though acknowledges that it is near the borderline). Must we deny that Freddy's verdict is knowledgeable? After all, it is, we may suppose, the verdict of someone who gives every indication otherwise of a normal competence in the concept, has normal vision and is judging in good conditions – and judging in a way consistent with his judgement of the same shade on other occasions. It is harsh to say he doesn't know.²¹ And if it is at least indeterminate whether Freddy knows, then we do not know VE.

On the other hand, if we take it that the EC-conditionals *are* known to hold good for surveyable predications of *F*, must we not also accept the strange claim that VE is known to be false for such cases and hence that each element in a soritical series for *F* allows in principle of a knowledgeable verdict about its *F*-ness? It's not clear. There is a double negation elimination step in the drawing of that conclusion whose legitimacy might be viewed as *sub judice* in the present dialectical context. Rather than take a stand on the

¹⁹ Schiffer (2016), Rosenkranz (2003), Williamson (1996b).

²⁰ Wright (2003a).

²¹ Some will no doubt say that Freddy has a different concept. But that seems merely ad hoc. What does the difference consist in? Why not say instead that he is steadier than we are in his judgements involving a concept we share?

matter, it may seem that prudence dictates, pro tem, that we reserve judgement on both VE and EC, committing to neither.

Prudence, though, comes at a cost. Unfortunately for the would-be intuitionist, that agnostic attitude requires that we must also be agnostic about the basic revisionary argument. If, in our current state of philosophical information, the strongest relevant claim we can justifiably make about the EC-conditionals is that it is epistemically possible that they hold good for surveyable predications of ‘yellow’, ‘bald’, etc., then, supposing we accept the validity of the law of excluded middle, we can validly reason only to the epistemic possibility that D above holds good, i.e. that it is *epistemically possible* that, for each P in the relevant class of statements, it is feasible to know P or it is feasible to know not- P . But that double-modalised conclusion doesn’t look uncomfortable – or anyway, not uncomfortable enough to put pressure on the acceptance of excluded middle. In particular, if it is epistemically possible that Steady Freddy indeed knows, then for each P in the relevant range, there epistemically possibly could be a steady subject who knowledgeably judges that it is true (or that it is false.)

The revisionary import of the basic revisionary argument requires more than that the EC-conditionals are epistemically possibly correct.

5.7 Addressing Constraint 1: A Different Tack – Knowledge-Theoretic Semantics

So what now? Well, a suspension, perhaps temporary, of confidence in the basic revisionary argument in this context need not surrender all prospect of a strong motivation for an intuitionistic approach to the logic of vague discourse. The basic revisionary argument attempts to garner the desired result without any particular assumptions about semantics. Let us therefore now instead consider directly what might be the most desirable shape for a semantics to take that is to be adequate for a language – a *minimally sufficient soritical language for F* – that has just enough resources to run instances for a particular vague predicate F of both the Tolerance and No Sharp Boundaries Paradox. Such a language thus contains the predicate F , a finite repertoire of names, one for each member of a suitable soritical series, brackets subject to the normal conventions, and the standard connectives and quantifiers of first-order logic.

Let L be such a language. Since we wish to avoid any commitment to the idea that when F is applied to an object that is surveyable for it, the result can take a truth-value beyond our ken, we have no interest in any semantic theory for L that works with an evidentially unconstrained notion of truth. But nor, since we are now (even if temporarily) agnostic about EC (and therefore also about VE), should such a semantics work instead with a verificationist notion of truth. It follows that we should not choose a truth-theoretic semantics at all.²² But then what? Well, what any competent practitioner of L has to master

²² I am not here assuming that merely to give a truth-theoretic semantics for some region of discourse must involve explicit commitment to one horn or the other of this alternative. But the question may legitimately be pressed, and the point I am making in the text is that we cannot answer unless at the cost of surrender of agnosticism about EC. Better, therefore, not to

are the conditions under which its statements may be regarded as known or not. We may therefore pursue a semantic theory that targets such conditions directly, in a spirit of aiming at a correct description of what we are in a position to regard as knowledgeable linguistic practice. It will be for the critic to make the case, if there is a case to make, that we thereby misdescribe the practice we actually have.

How to make a start? We don't have much to go on. What is solid to begin with is only that there is a range of polar cases where there is no doubt that Fx may be known, a range of polar cases where there is no doubt that $\sim Fx$ may be known and a range of cases that manifest the uncertainty and variability of judgement that our governing deflationism regards as constitutive of vagueness. But consider the following controversial principle (CP):

All the *knowable* statements in L are knowable by means of knowing the truth-values of atomic predications – (which we are assured of being able to do only in polar cases.)

According to CP, any of the molecular statements of L , can be known, if they can be known at all, by knowing some of L 's atomic statements. So the semantic clauses for the connectives and quantifiers by means of which any particular molecular statement is constructed ought – if that statement is to be reckoned knowable – to reflect an upwards path, as it were, whereby the acquisition of such knowledge might proceed. If we accept CP, we will be looking therefore for a semantics that recursively explains conditions of knowledge for the molecular statements of L in terms of those of their constituents.²³

Presumably, we are not going to want to accept CP. 'Controversial' somewhat flatters the principle. We will surely want to admit a range of exceptions, cases where a molecular statement plausibly holds good even when its constituents are borderline. Some, for instance, will be general statements that are arguably analytic of the specific vague predicate concerned, such as 'Everything red is coloured'; others may be nomologically grounded in the property concerned, such as maybe 'All heaps are broadest in the base'. A more interesting class of exceptions are what Kit Fine once characterised in terms of the notion of *penumbral connection*.²⁴ They will concern vague predicates in general. Epistemicists will regard some instances of the law of excluded middle as coming into this category. We will not follow them in that, but we should want to allow, for example, that no matter what F may be, all instances of the law of non-contradiction are knowable as, with respect to the kind of series we are concerned with, are all *monotonicity conditionals*, that is, statements of the form

$$Fx' \rightarrow Fx,$$

invite the question. However, there is more to say about the motivation for the style of semantics about to be proposed. I'll return to the matter at the end.

²³ For ease of formulation, I here count the instances of a quantified statement as among its 'constituents'.

²⁴ Fine (1975).

notwithstanding whether x is borderline for F . The same will hold for the corresponding generalisations:

$$(\forall x) \sim (Fx \ \& \ \sim Fx), (\forall x)(Fx' \rightarrow Fx), (\forall x)(\sim Fx \rightarrow \sim Fx'), \dots$$

To be sure, that such claims are knowable is not uncontroversial. It is a familiar feature of many-valued treatments of vagueness that such principles as these are sometimes parsed as indeterminate – when, for instance, indeterminacy in a conjunct is treated as depriving a conjunction of determinate truth, or a conditional with an indeterminate antecedent and consequent is regarded as thereby indeterminate. We are not here taking a stand on the question whether such treatments are appropriate when one accepts their governing assumption, viz. that being a borderline case is a kind of *alethic* status, contrasted with both truth and falsity. But we are rejecting the governing assumption. And when instead borderline-case status is viewed as a cognitive status, as on *our* governing assumption, there is no evident reason to demur at the suggestion that principles of penumbral connection can be known. We can know of structural constraints that knowledge, were it but attainable, of the truth-values of a range of statements would have to satisfy without having any guarantee that we can get to know those truth-values.

These considerations suggest we pursue a theory of knowledge for L that has CP as a motivating base but includes a range of permitted exceptions to it. The theory will incorporate a knowledge-conditional semantics for L and a logic based upon it, but may also contain additional, primitive axioms of penumbral connection and perhaps other axioms analytic of or otherwise somehow guaranteed for a particular choice for F . The semantics will comprise recursive clauses that determine, for each of the quantifiers and connectives of L , the conditions that are necessary and sufficient for knowledge of L -statements in which that operator is the principal operator on the basis of the knowledge-conditions of its constituents.

The natural approach will be something in the spirit of the Brouwer-Heyting-Kolmogorov (BHK) interpretation of intuitionist logic,²⁵ which, as is familiar, proceeds in proof-theoretic rather than truth-theoretic terms. There is, however, an important point about the BHK interpretation that we need to flag before moving to propose specific clauses for the theory for L that we seek. In logic and mathematics, or so one might plausibly hold, all knowledge (other than of axioms) is conferred by, and only by, proof. So it can look as though BHK-style semantics is already nothing other than a local version of knowledge-conditional semantics. So it is, but expressing matters that way may encourage an oversight. While proofs in logic and mathematics confer knowledge of what they prove, that is not all they do. They also vouchsafe knowledge of what is proved *as* knowledge. Someone who comprehendingly works through a mathematical proof that P learns not merely that P is true but also – assuming their grasp of the concept of knowledge, etc. – that P may now be taken to be part of their knowledge. They establish a right to include P

²⁵ See e.g. Troelstra (2011, section 5.2).

as part of what they may legitimately claim to know. Say that knowledge is *certified* when accomplished in a fashion that legitimises that claim: accomplished in such a way that a fully epistemically responsible, sufficiently conceptually savvy epistemic agent will be aware that they have added to their knowledge. The clauses to follow are to be understood in terms of knowledge that is certifiable – *c-knowledge* – in this sense.²⁶

Adapting BHK-style clauses in a natural way, we may accordingly propose:

- ' $A \ \& \ B$ ' is knowable just if it's knowable that ' A ' is knowable and that ' B ' is knowable.
- ' $A \ \vee \ B$ ' is knowable just if it's knowable that either ' A ' is knowable or ' B ' is knowable.
- ' $(\forall x)Ax$ ' is knowable just if it's knowable that for any object in the soritical domain and term, ' a ', known to denote that object, ' Aa ' is knowable.²⁷
- ' $(\exists x)Ax$ ' is knowable just if it's knowable that, for some object in the soritical domain and term, ' a ', known to denote that object, ' Aa ' is knowable.

(What about the conditional? Actually, we don't strictly need a treatment of the conditional for the present purposes.²⁸ And this is fortunate, since the natural proposal:

' $A \rightarrow B$ ' is knowable just if it's knowable that if ' A ' is knowable, ' B ' is knowable, raises an awkwardness which I will explain below.)²⁹

We can now assert the following Thesis (verification is left to the reader):³⁰

Where validity is taken as *c-knowability-preservation*, and *c-knowledge* is taken to be *factive* and *closed over c-knowable logical consequence*, the clauses above justify rules of deduction for the listed operators coinciding with the common ground for those operators – the standard rules for conjunction introduction and elimination, disjunction introduction and elimination, universal generalisation and instantiation, and existential generalisation and instantiation – recognised by both classical and intuitionist first-order logic.

What about negation? An adaptation of the BHK-style clause along the above lines would run:

' $\sim A$ ' is knowable just if it is knowable that ' A ' is not knowable.

²⁶ It is a consequence of some kinds of knowledge-externalism that not all knowledge need be *c-knowledge*. But the externalist will presumably grant that knowledge often is *c-knowledge*, since it is not supposed to be a consequence of externalism that our claims to knowledge are mostly imponderable without further investigation. The crucial assumption I am making in what follows is that knowledge achieved by canonical means – typically, casual observation – of clear cases of the 'usual suspects' will be *c-knowledge*.

²⁷ Recall that *L* will contain a known name for every element of the soritical domain.

²⁸ When the major premise for the Tolerance Paradox is formulated, as standardly, as a universally quantified conditional (rather than e.g. as involving a binary universal quantifier) then the paradox does of course depend on the unrestricted use of *modus ponens*. But the intuitionist resolution of the paradox to be proposed will pick no quarrel with that, and is thus neutral on the semantics of the conditional to that extent.

²⁹ See fn. 31.

³⁰ 'And how', the dear reader may ask, 'am I supposed to do that when you have nominated no specific logic for the meta-language – here English! – in which I am supposed to run through the relevant reflections?' *Touché*. But the meta-reasoning concerned will require, besides the noted properties of *c-knowledge*, no more than the rules of inference for 'and', 'or', 'any', 'some' and 'if' which constitute common ground between classical and intuitionist first-order logic.

But that, obviously, will introduce calamity into any account that accepts VE. Our official stance at this point is one of agnosticism towards VE, but it would be good to have the resource of a treatment of the paradoxes that would be robust under the finding that VE was after all philosophically mandated. In any case, and perhaps more telling, BHK-style negation has always been open to the intuitive complaint that it provides a licence to convert grounds for thinking we are doomed to ignorance on some matter into grounds for denial and thus distorts negation as intuitively understood.

There is, however, a natural and much more intuitive replacement:

‘ $\sim A$ ’ is knowable just if some ‘ B ’ is knowable that is knowably incompatible with ‘ A ’,³¹

or more generally

‘ $\sim A$ ’ is knowable just if some one or more propositions are knowable that are conjointly knowably incompatible with ‘ A ’.

There is no space here to undertake a proper exploration of the philosophical credentials of this proposal. Still, the reader may find it intuitively plausible that mastery of negation, at least at the level of atomic statements, is preceded in the order of understanding by mastery of which of them exclude which others: being not yellow, for example, is, among coloured things, initially understood as the having of some colour that rules out being yellow. The above proposal reflects the thought that we may take incompatibility among atomic statements as epistemically primitive. Matters change, of course, once molecular statements enter the mix. For molecular statements, incompatibility will, conversely, sometimes be recognisable only by recognising that one or more of them entail the negation of something entailed by the other. That is,

³¹ The reader should note that there is a question, drawn to my attention by Timothy Williamson, whether we may stably combine this proposed knowledge-theoretic clause for negation with the knowledge-theoretic clause for the conditional flagged earlier:

‘ $A \rightarrow B$ ’ is knowable just if it’s knowable that if A is knowable, B is knowable.

For suppose VE is accepted and A is such that

- (i) It is knowable that A is borderline.
- (ii) Then it is knowable that A is not knowable (by VE).
- (iii) So it’s knowable that if A is knowable, then B is knowable (by substitution in $\sim A \Rightarrow (A \rightarrow B)$ (*ex falso quodlibet*) and closure of knowledge across knowable entailment).
- (iv) So it’s knowable that if A , then [take some arbitrary contradiction for B] (from iii, by the knowledge-theoretic clause for the conditional).
- (v) So it’s knowable that $\sim A$ (by the proposed clause for negation, letting B be: if A , then [contradiction], and presuming that to be knowably incompatible with A).

So the proposed clause degrades after all into the BHK-style knowledge-theoretic clause for negation:

‘ $\sim A$ ’ is knowable just if it is knowable that ‘ A ’ is not knowable,

which is what we were trying to improve on. True, the argument as presented depends on VE, which we have not endorsed. But it will run for any ‘ A ’ that is knowably unknowable. Maybe there are no such statements formulable in a minimally sufficient soritical language. Maybe one should look askance at *ex falso quodlibet*. Still, a concern is raised that will need disinfection in a fully satisfactory general treatment. I will not pursue the matter here.

A pair of (sets of) propositions are knowably mutually incompatible if there is some proposition ‘A’ such that the one knowably entails ‘A’ and the other knowably entails ‘~A’.

If we now, for convenience, avail ourselves of a dedicated constant, ‘⊥’, to express the situation when a set of propositions, X, incorporates each of some pair of knowably incompatible propositions, thus:

$$X \Rightarrow \perp,$$

then the first of the displayed clauses above mandates the following negation introduction rule:

$$(\sim\text{Intro}) \frac{X \cup \{A\} \Rightarrow \perp}{X \Rightarrow \sim A}$$

while the second displayed clause mandates the following negation elimination rule:

$$(\sim\text{Elim}) \frac{X \Rightarrow \sim A \quad Y \Rightarrow A}{X \cup Y \Rightarrow \perp}$$

That is, intuitively, if a set of propositions entails the negation of some proposition, then adding to it any set of propositions that entail that proposition will result in incompatibility.

Whatever deep justification these proposals may be open to, it will be enough for present purposes if they seem plausibly knowability-preservative in the light of the reader’s intuitive understanding of negation. Their most immediately significant consequence is that they allow us to justify intuitionist *reductio* as a derived rule.³² Given the Thesis flagged above, we thus have all the rules (&I, ∃I, *reductio*) needed to run both the No Sharp Boundaries Paradox and – assuming no question is raised about *modus ponens* – the Tolerance Paradox as well. The upshot, in the presence of assumed knowledge of the polar assumptions, is the following important corollary:

Corollary: When the quantifiers and connectives are understood as above, *there is no option* but to regard the negations of the major premises of the No Sharp Boundaries Sorites as known.

³² At least they do so if we may assume the Cut rule. Intuitionist *reductio* may be represented as the pattern:

$$\frac{X \cup A \Rightarrow B \quad Y \Rightarrow \sim A}{X \cup Y \Rightarrow \sim A}$$

Suppose we have an instance of right-hand premise. From that and $B \Rightarrow B$ we may obtain by ~Elim:

$$Y \cup \{B\} \Rightarrow \perp$$

From that and the left-hand premise we have, by Cut

$$X \cup Y \cup \{A\} \Rightarrow \perp$$

So by ~Intro, we have

$$X \cup Y \Rightarrow \sim A$$

5.8 Addressing Constraint 1 (Cont.): The Payoff

Constraint 1 requires that we explain how and why the *reductio* of the major premises accomplished by the paradoxical reasoning fails to justify the unpalatable existential. This is now straightforward. By the clause for ‘ \exists ’, the knowability of the unpalatable existential requires that for some object in the soritical domain and term, ‘ a ’, known to denote that object, ‘ $Fa \ \& \ \sim Fa$ ’ is knowable – requires, in short, the knowability of a witness to a sharp cut-off. We neither have nor have any reason to think we can obtain that knowledge: no ‘ a ’ denoting any clear case furnishes such a knowable witness. And we have absolutely no reason, either given by the Sorites reasoning itself or otherwise, to think that such a witness may be knowledgeably identified in the borderline area. Since, by the Corollary emphasised at the conclusion of the preceding section, we do know the negations of each of NSB and TP, the respective classical inferences from the negations of NSB and TP to the unpalatable existential fail to guarantee knowability and are thus are invalid in the present knowledge-theoretic setting.

So there is the needed daylight. Constraint 1 is met and the discomfort involved in regarding the soritical reasoning simply as a *reductio* of its major premise is thus relieved.

5.9 Addressing Constraint 2

At least, it is relieved if, as required by Constraint 2, we can neutralise the persistent temptation to regard the major premises for the paradoxes as true. The epistemicist – and indeed almost all theorists of this topic³³ – also share this obligation, of course, so here we, most of us, can march in step. There are a number of sources for the temptation. I’ll touch on four.

(i) *Projective error*

The core attraction of NSB is, naturally, simply the other face of the unpalatability of the Unpalatable Existential. And that in turn springs from our inclination to accept that NSB is simply a statement of what it is for F to be vague in the series in question. On our overarching conception of what vagueness is, this is a tragic mistake. It is, indeed, the pivotal mistake, ‘the decisive step in the conjuring trick’ that our intuitive thinking plays on us here. For F to be vague is for it to have borderline cases, but its possession of borderline cases is, according to the overarching deflationary conception of vagueness here proposed, a matter of our propensity to certain dysfunctional patterns of classification outside the polar regions.³⁴ F ’s being vague is thus a fact *about us*, not about the patterns that may or may not be exhibited by the F s and the non- F s in a Sorites series. Nothing follows from

³³ The exceptions are those theorists who prefer to look askance at the underlying logic of the paradox; for instance at the assumption of the transitivity of logical consequence in this setting (Zardini, see Chapter 9 in this volume) or at intuitionist *reductio* (Fine 1975).

³⁴ This much is common ground with epistemicism as I understand it. The difference is that we reject the further step of postulating a sharply bounded property, our inability to keep track of whose extension explains the dysfunctionality.

its vagueness about thresholds, or the lack of them, in a Sorites series, or indeed about the details of its extension at all.

The diagnosis of projective error chimes nicely with Constraint 3: the overarching conception of borderline case vagueness we are working under is invoked to undergird the proposed means of satisfying Constraint 2.

(ii) *Inflated normativity*

However, there are other kinds of seductive untruth at work in conjuring the attraction that the major premises exert. One such involves an implicit inflation of the legitimate sense in which competent practice with the usual suspects is constrained by *rule*, and is a crucial factor in the allure of tolerance premises. An example is the general thought that the rules for the use of any of the usual suspects that can be justifiably applied or denied purely on the basis of (varying degrees of casual) *observation*, must mandate that elements in the soritical domain between which there is no relevant (casually) observable difference should be described alike. (For how otherwise could mere observation enable us to follow the rules?) In fact, none of the expressions with which we are concerned is governed by rules that mandate any such thing. But the illusion that they are – indeed must be – so governed has deep sources. As announced earlier, I must forbear to go further into these matters here.³⁵

(iii) *An operator shift?*

Both the foregoing, though ultimately misguided, are nevertheless respectable reasons for our inclination to accept the major premises, involving subtle philosophical mistakes. I am not completely confident that some of us, over the last four decades of debate of these paradoxes, may not have fallen prey to a less respectable reason. (I am sure no present reader would be guilty of this.) There is a fallacious transition available in this context of a kind that we know it is easy to slip into: an operator shift fallacy. The transition concerned is that from

Nothing in the meaning of F (the way we understand it) mandates a discrimination between adjacent elements in the soritical series,

to

The meaning of F (the way we understand it) mandates that there is to be no discrimination between adjacent elements in the soritical series, so that ' $Fx \ \& \ \sim Fx$ ' is everywhere false.

(iv) *Irrelevant truths*

Finally there are a number of truths in the vicinity that may tempt one to accept NSB but which, on the present deflationary conception of vagueness, adjoined with liberalism,

³⁵ Wright (1975) rehearses what is still the best case known to me for thinking otherwise, and points an accusatory finger at the implicit inflation of normative constraint; the inflation is further explained and debunked in Wright (2007).

simply have no bearing on it. It is true, for example, that no clear cases bear witness to the unpalatable existential, that nobody could justify claiming to have identified a witness in the borderline area and that we (normal speakers) have no conception of what it would be like even to have the impression that we had identified a witness. But these all merely reinforce the impression that UE, the unpalatable existential, is nothing we can justify. The mere possibility of a coherent epistemicism should teach us, if nothing else, at least that such considerations do not parlay into good reasons for its *denial*.

The temptation though dies hard. ‘Granted’, it may be said, ‘that if we are epistemicists, considerations like the above provide no good reason for denial. But what if we are not epistemicists? What if our attitude is, as the governing conception of vagueness that you are proposing itself involves, that there are here no facts behind the scenes: that our best practice exhausts the relevant facts – that “nothing is hidden”? If there are no truth-makers for predications of F and not- F save aspects of our best practice, then – given that our best practice determines no sharp cut-off for F in a suitable series, must we not conclude that there is none? And then is not some version – employing some suitably “wide” notion of negation – of NSB going to be forced on us?’

The question is in effect: how can we avoid treating a fact about us and our judgemental limitations as a fact about the properties of the elements of the series unless we are prepared, with the epistemicist, to invoke some form of transcendent fact? I reply that, while our judgemental reactions no doubt are caused by and reflect properties of the elements of the series, it is a further, unwarranted step to draw conclusions from them about the extension of F . The transition from

Our collective best practice doesn’t converge on a sharp boundary for F in the series

to

There is no sharp boundary for F in the series

is still a non sequitur even if we do not admit any practice-transcendent truth-makers for the statements in question. What does follow is at most a conclusion about indeterminacy – that something has not been settled by a convergence in our collective practice. But indeterminacy, conceived as proposed by the intuitionist, is not an alethic status, so *not something that excludes truth*.

Is this too much to swallow? Let the critic have another go: ‘Suppose it agreed’, she may say, ‘that there are just two kinds of ways in which an instance of Fx & $\sim Fx'$ can hold true in the series. One is the epistemicist way. The other is as grounded in the linguistic practice of competent judges in good conditions. Suppose we reject epistemicism. Then surely we are forced to accept the conditional:

If, for no element in the series, is there sufficient agreement among competent judges in good conditions on the truth of the relevant instance of Fx & $\sim Fx'$, then no such statement is true?

So then, since there is no foreseeable such agreement, each such statement is untrue; and now it must be possible to run an NSB Sorites in terms of a suitably wide negation.'

What obstructs this train of thought, as the reader may anticipate, is liberalism. If we accept liberalism about verdicts in the borderline area, we must reject the displayed conditional anyway, even without epistemicism as a background assumption. Liberalism requires that we not insist on available convergence about an atomic statement among competent judges as a necessary condition for its truth. Why should it be any different for molecular statements in general and an instance of $Fx \ \& \ \sim Fx'$ in particular?

'Not so fast', the critic may continue. 'If we are to leave open the possibility of an instance of $Fx \ \& \ \sim Fx'$ holding true, and if this possibility is not to be understood as the epistemicist understands it – as a matter of a cut-off in the extension of a property that is the semantic value of F but of whose nature we are not fully aware – and if, moreover, the possibility is not to be understood, either, as realised by a convergence in our best practice on each conjunct, then what is it a possibility of? What other kind of state of affairs, if it obtained, could conceivably be a truth-maker for an instance of $Fx \ \& \ \sim Fx'$ '?

The critic is assuming that there can be no truth without a truth-maker. But let us not question that and consider how an intuitionist might answer her question directly. Let ' Fa ' be a non-polar statement. Suppose that a is the last element in the Sorites series for which a competent judge in good conditions – Freddy again – returns a steady positive, if suitably nuanced, verdict. Liberalism requires that we not discount Freddy's verdict about a . But suppose also that a' is the first element in the series for which Teddy, Freddy's epistemic peer, returns a steady negative, if suitably nuanced, verdict. Liberalism requires that we not discount Teddy's verdict either. Should we not then be liberal towards their conjunction? The displayed conditional, however, would force us to dismiss the conjunction of Freddy's and Teddy's respective verdicts as untrue, for that conjunction elicits nobody's assent, however competent, however good the conditions.

The critic may be unpersuaded. 'One can perfectly reasonably be liberal about a pair of verdicts individually but illiberal about their conjunction? Change the example. What if Teddy and Freddy were to steadily disagree about whether Fa ? Now there is no option of regarding both as right – yet that does not preclude our taking a liberal view of each verdict on its own. So why should liberalism about Freddy's and Teddy's respective verdicts either side of the putative cut-off provide any leverage towards liberalism about them taken together?'

I reply that such leverage is the default: that liberalism about any pair of judgements individually should extend to their conjunction except in cases where there is antecedent reason to recognise tension – for example, flat contradiction! – between the judgements concerned. Unless, therefore, one is *independently* inclined to see the truth of ' Fa ' as in tension with the truth of ' $\sim Fa$ ', there is no reason to look askance at the conjunction of Freddy's and Teddy's respective steady verdicts. If, however, you consider that you do have good reason to be independently so inclined, you will presumably be independently inclined to accept NSB. The resurgent paradox will then be a deserved nemesis.

5.10 Addressing Constraint 3

The third constraint we imposed on an intuitionistic treatment of the Sorites was the requirement that the first two constraints – explaining how and why there can be a deductive gap between the negation of the major premise and the unpalatable existential, and explaining the spurious plausibility of the different forms of major premises – be met in a way that is informed by an overarching conception of what vagueness consists in. Have we done this?

It is arguable that the second constraint is not really motivated in the case of the Tolerance Paradox. To be sure, the vagueness of a predicate, deflationarily conceived, is nothing that should suggest that it be tolerant. However, the principal motivations to regard e.g. the usual suspects as tolerant have, as remarked, little to do with their vagueness per se and need a separate treatment, not embarked on here. On the other hand, the diagnosis of projective error as responsible for the thought that an NSB premise just states what it is for F to be vague in a relevant soritical series draws heavily and specifically on the deflationary conception of vagueness that I have represented as the heartbeat of an intuitionistic approach.

But what about the first constraint – explaining the deductive gap? Assume that the knowledge-theoretic semantics offered performs as advertised. We have to acknowledge that a semantic theory of this kind might be proposed, for certain purposes, for almost any factual discourse. So the question becomes: what, if anything, is it about vagueness as deflationarily conceived that makes such a semantics appropriate for vague discourse specifically?

Recall that it is essential to our deflationism not merely to regard certain judgemental patterns among competent judges as constitutive of an expression's vagueness but to reject the demand for explanation of these patterns in terms of underlying semantic phenomena – for instance, sharply bounded but imperfectly understood semantic values, incomplete (or conflicting) semantic rules or the worldly side of things being such as to confer truth-statuses other than truth and falsity. That precludes any semantic theory that works with a bivalent notion of truth, truth-value gaps or postulates any kind of third truth-status. Admittedly, the possibility is left open of working with a verificationist truth-conditional theory, as would be mandated by EC. However, no reason is evident why the knowledge-theoretic style of semantics proposed could not amount to one way of implementing the semantic import of EC, nor hence why all the crucial parts of the treatment of the paradoxes proposed could not survive were we to quash any reservations about EC. So we have not closed that particular road by going about things the way I have here. But nor have we committed to travelling it.³⁶

³⁶ Versions of this material were presented at the Staff Research Seminar at Stirling, the Philosophy of Maths seminar at Oxford, an Arché seminar at St Andrews in the autumn of 2016, and at colloquia at Brown University and the University of Connecticut in the spring of 2017. I am grateful to all who participated in these meetings for useful feedback, and to Ian Rumfitt and Josh Schechter for helpful additional discussion. Special thanks to the editors of this volume for very searching critical comments that have led to many improvements.