

# Deep Ignorance, Brute Supervenience, and the Problem of the Many

Terry Horgan

Timothy Williamson holds that vagueness, properly understood, is an epistemic phenomenon: vague terms and concepts have precise boundaries of correct applicability and non-applicability, but humans are necessarily ignorant of the locations of these boundaries. Williamson argues that our ignorance of boundary locations, and our inability to imagine recognizing a transition as a transition, are predictable on the basis of the epistemic theory's account of how vagueness works, and hence that these facts do not count against the theory. The sort of explanation he offers, as I understand it, appeals (i) to the claim that the specific precise boundaries for any vague term are sensitively dependent on the details of the overall pattern of usage upon which applicability of that term supervenes, and (ii) to the fact that we are not able either to ascertain or to cognitively survey all relevant aspects of such a use-pattern.<sup>1</sup> Concerning the unintuitiveness of epistemicism, he says:

---

<sup>1</sup>Williamson evidently thinks that there are additional cognitive limitations, over and above those mentioned in (ii) but similar in kind, that would prevent us

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. (p. 221)

I think one should grant that we cannot ascertain and cognitively survey the use-patterns on which our concepts supervene, and that hence, even if the concepts have precise boundaries, we would not be able to know where they are or be able to imagine recognizing a boundary transition as a transition. However, I also think that there is a deeper, more fundamental, kind of unimagability phenomenon than our inability to imagine recognizing a transition *qua* transition. Furthermore, I think that this deeper kind of unimagability grounds a powerful and highly persuasive argument against epistemicism. I suspect that the argument in question, or something very much like it, is what lies behind the fact that most of us find epistemicism too incredible to believe.

I will briefly set out that argument here. I will also briefly consider some avenues that an epistemicist might try as ways to block the argument. Each of those avenues, I will argue, involves theoretical commitments of a radical and implausible kind.

Consider a snowmass on a mountainside, shortly prior to the onset of an avalanche. In principle, we may suppose, the precise moment of onset of the avalanche is predictable from the intermolecular state of the snowmass, on the basis of known laws of physics. But in practice, producing such a prediction on such a basis is completely out of the question. We couldn't possibly obtain a complete description of the total intermolecular state of a large snowmass, or cognitively survey such a gargantuan description, or perform all the billions upon trillions of calculations necessary to precisely predict the moment of avalanche-onset on the basis of that total-state description plus the

---

from being able to recognize a transition as a transition even if, *per impossibile*, we could ascertain and cognitively survey an entire use-pattern. He writes, "The same dialectic applies at the metalinguistic level. . . . 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" (p. 221–2). Likewise, in his *Vagueness* (Routledge, 1994), he writes, "You have no way of making your use of a concept on a particular occasion perfectly sensitive to your overall pattern of use, for you have no way of surveying that pattern in all its details. . . . Even if you did know all the details of the pattern (which you could not), you would still be ignorant of the manner in which they determined the content of the concept" (pp. 231–32). He does not say explicitly what the additional cognitive limitations consist in.

laws of physics.<sup>2</sup> (These things would remain impossible in practice even with the aid of the best measuring instruments and computers; the task is utterly intractable.) But the reasons why this is so are mundane. There is no deep mystery about avalanches, because we do know the fundamental principles that comprise the *explanatory basis* for avalanche-phenomena —viz., the relevant laws of physics.

The case of the snowmass is strikingly disanalogous to cases involving vague terms like ‘heap’. Although it is true enough that humans cannot ascertain and cognitively survey the overall pattern of use of the term ‘heap’ on which the term’s meaning supervenes, and although it may well be true too that this fact by itself suffices to explain why we would not be able to imagine recognizing a sharp transition from heaphood to non-heaphood even if such transitions exist, there is something much more fundamental that we also cannot imagine. We cannot imagine *why* it should be the case that any particular candidate for precise boundaries for the term ‘heap’, as opposed to numerous other candidates, should be the unique boundaries that actually supervene on the total use-pattern governing this term. That is, we cannot imagine what could possibly constitute the explanatory basis for the putative fact (call it a “precise supervenience fact”) that such-and-such specific precise boundaries supervene on so-and-so overall pattern of usage. In the case an avalanche, we know what constitutes an explanatory basis of avalanche-onset —viz., the relevant laws of physics. But in the case of vague terms like ‘heap’, we cannot even conceive what kinds of facts, principles, or laws could possibly play an analogous explanatory role vis-a-vis putative precise supervenience facts. Even *given* some specific overall use-pattern as the total supervenience base for the meaning of ‘heap’, it seems there still would be numerous equally good candi-

---

<sup>2</sup>Even if, *per impossibile*, we could cognitively survey a complete description of the total intermolecular state of a snowmass, we still would not be able to apply the laws of physics to this gargantuan description in order to generate a predication of an avalanche; there would be far too many calculations to perform. Likewise, even if, *per impossibile*, we could ascertain and cognitively survey the total use-pattern on which the correct application of the word ‘heap’ supervenes, and even if we knew certain explanatory principles that underwrite precise supervenience facts, we still might very well be unable to apply those principles in order to ascertain the specific precise boundaries of ‘heap’ —or in order to ascertain a transition between patterns of use in virtue of which the word correctly applies to a given pile and patterns of use in virtue of which it does not apply. Thus, on one natural construal, the limitation to which Williamson adverts (see note 1) as a metalinguistic counterpart of our inability to cognitively survey the relevant kinds of use-patterns is this limitation in our deductive/calculational capacities.

dates for precise boundaries, and there would be no reason whatever why one of these candidates should win out over the others.

Adapting Peter Unger's terminology, let me call this deeper problem of unimagability *the Problem of the Many*. It is an instance of a class of problems, identified by Unger under this rubric, exhibiting a common form —viz., numerous equally eligible precise candidates for identity with some item (an object, a meaning, a concept, etc).<sup>3</sup> The specific Problem of the Many I am stressing here is this: among the many equally eligible-looking candidate boundaries that might potentially be supervenient on any given use-pattern, we cannot conceive or imagine any explanatory basis that would "break symmetry" among the various candidates and would single out some unique candidate over against all of its competitors.

The unimagability featured in the Problem of the Many is evidently orthogonal to the fact that we cannot ascertain and cognitively survey the overall pattern of usage on which a concept or a meaning supervenes. For again, the problem is that we cannot imagine what could possibly explain why this pattern, whatever its specific details might be, should be the supervenience base for any *single* precise boundaries over against numerous other candidate boundaries. (In the case of the snowmass, on the other hand, although here too we cannot ascertain and cognitively survey the relevant intermolecular initial conditions that determine the moment when an avalanche will commence, we *do* know what explains the connection between the current state and the subsequent avalanche —viz., certain physical laws.)

The Problem of the Many, I would maintain, is the unimagability problem that leads so many of us (myself included) to believe that vague terms and concepts do not, and cannot, have precise boundaries. The reasoning can be reconstructed as follows:

- (1) We cannot conceive or imagine an explanatory basis for precise supervenience facts involving vague terms and concepts —that is, for facts of the form, "Such-and-such precise boundaries supervene on so-and-so overall pattern of usage".
- (2) So there is no such explanatory basis.
- (3) So there are no such supervenience facts.
- (4) So there are no precise boundaries for vague terms and concepts.

---

<sup>3</sup>Peter Unger, "The Problem of the Many", *Midwest Studies in Philosophy* 5, 1980, 411–67.

Although the successive steps in this argument are nondeductive, they are each extremely plausible.

How might an epistemicist try to block this reasoning? The initial premise cannot be credibly denied, I take it. Nor is it credible to deny the inference from (3) to (4); and Williamson, certainly, commits himself explicitly to the claim that meanings and concepts supervene on other facts —roughly, facts about use. This leaves two principal options: denying the inference from (1) to (2), or denying the inference from (2) to (3).

One way of rejecting the inference from (1) to (2) would be to claim that there is indeed some explanatory basis that humans can understand, even though as yet we cannot conceive such a thing. This claim is not unintelligible; it is arguable, after all, that before Einstein nobody could conceive of the relativity of simultaneity, or of curved spacetime. But the problem is that there is presently no more than a *bare epistemic possibility* that some future Einstein could discover the explanatory principles that we presently cannot even conceive. There is no good reason to believe this; and there is excellent reason to disbelieve it, viz., the current inconceivability of such principles.

Another way of rejecting the move from (1) to (2) would be to make these two claims: first, that there is indeed an explanatory basis for precise supervenience facts; but second, that humans are inherently incapable of grasping the relevant explanatory principles. I will call this the hypothesis of *deep ignorance*, since it posits certain unknowable, humanly incomprehensible, explanatory facts. Such a view would be the analog, for vagueness, of Colin McGinn's position concerning the supervenience connections between neuro-chemical properties and phenomenal mental properties ("qualia"). McGinn holds that although there is some naturalistic explanation for these physical/phenomenal supervenience relations, human beings are (as he puts it) *cognitively closed* to this explanation.<sup>4</sup>

Needless to say, a hypothesis of deep ignorance is enormously radical. It is the kind of philosophical position one should embrace, if at all, only as a theoretical last resort. Many in the philosophy of mind find McGinn's deep-ignorance hypothesis too radical to be credible, but at least there is this to be said for it: the "what-it's-like" aspects of phenomenal consciousness are undeniably a part of the pre-theoretic phenomena that need accommodating within an overall philosophical position about mind and brain. The pre-theoretic salience of these subjective, experiential, aspects of phenomenal

---

<sup>4</sup> *The Problem of Consciousness*, Blackwell, 1991, Chapters 1-4.

consciousness is crucially important in generating epistemic pressure toward a view like McGinn's. On one hand, we cannot credibly deny the phenomena in question; but on the other hand, it is very difficult for many philosophers (myself included) to conceive any way of explaining supervenience connections linking neuro-chemical properties to these subjective "what it's like" properties.

In the case of vagueness, however, there is not a comparable kind of epistemic pressure toward a deep-ignorance hypothesis. On the contrary, insofar as one considers relevant pre-theoretic phenomena, the situation is quite the reverse: common-sense intuition strongly suggests that vague terms and concepts do *not* have precise boundaries. Thus, not only is the idea of deep ignorance just as radical *vis-a-vis* vagueness as it is *vis-a-vis* phenomenal consciousness, but in the case of vagueness it flies in the face of the pre-theoretic data rather than being fueled by them. Whether or not one finds the idea credible for consciousness, in the case of vagueness it is apt to seem wildly implausible. It seems so to me.

Suppose an epistemicist leaves intact the inference from (1) to (2) in the above argument, and instead questions the inference from (2) to (3).<sup>5</sup> Here the claim would be that although there are precise supervenience facts —i.e., it is indeed the case that precise boundaries for vague terms and concepts supervene on total patterns of use— nevertheless such supervenience facts are metaphysically brute and *sui generis*, rather than being explainable. I will call this the hypothesis of *brute supervenience*.<sup>6</sup> The view could take various forms, depending on the strength of the modality involved. One variant, brute *metaphysical* supervenience, would claim that there are no possible worlds whatever that are just alike with respect to use-patterns but different with respect to precise boundaries for vague terms and concepts. Another variant, brute *nomological* supervenience, in effect would posit certain synchronic supervenience-laws as fundamental laws of nature alongside the basic laws of physics.<sup>7</sup>

---

<sup>5</sup>In the discussion of his paper during the SOFIA conference, Williamson expressed an inclination to take this tack.

<sup>6</sup>The metaphysically basic supervenience principles might be inter-level laws with general applicability —so that particular supervenience facts, linking specific use-patterns to specific boundaries, would be instantiations of these laws. But there would still be brute supervenience, albeit at the level of the general laws rather than at the level of particular instantiations.

<sup>7</sup>This would be an analog of David Chalmers' position concerning supervenience connections between neuro-chemical properties and phenomenal mental properties, in his *The Conscious Mind: In Search of a Fundamental Theory*, Oxford, 1996. Chalmers acknowledges that his view is radical, because it posits

On either variant, this is a very radical position. It entails, for instance, not only that the word 'heap' has the same perfectly precise boundaries in all possible worlds (or: in all nomologically possible worlds) in which 'heap' has the same use-pattern as in the actual world, but also that *there is no reason* why it is the case that in all these worlds, 'heap' has just these specific precise boundaries rather than any of various other candidate precise boundaries. Brute modal facts of this kind, if such there be, are incredibly peculiar; they are metaphysical surds, pockets of utter arbitrariness in the fundamental fabric of reality. Brute metaphysical supervenience is especially peculiar if, as seems plausible, virtually all other metaphysical necessities are grounded in non-tendentious *conceptual* truths.<sup>8</sup> But brute nomological supervenience is extremely queer too, not only because it unparsimoniously bloats the fundamental laws of nature beyond those of physics, but also because of the profound weirdness of fundamental laws that arbitrarily break ties among otherwise equally eligible candidates for precise boundaries.

It would be hard enough to believe in such surd modal facts even if precision of semantic/conceptual boundaries were part of the pre-theoretic phenomena of vagueness. Given that the pre-theoretic phenomena instead include the apparent absence of precise boundaries, the hypothesis of brute metaphysical supervenience for precise boundaries is apt to seem quite fantastic. It seems so to me.

In summary: There is an unimaginability phenomenon associated with vagueness that is deeper than our inability to imagine recognizing a precise semantic/conceptual transition-point in a sorites series: viz., our inability to imagine any explanatory basis for precise supervenience facts. This deeper phenomenon is the basis for a powerful-looking argument for the conclusion that vague terms and concepts do not have precise boundaries. Each of the two principal options open to the epistemicist for trying to block the argument—the hypothesis of deep ignorance, and the hypothesis of brute supervenience of precise semantic/conceptual boundaries—is radical,

---

non-physics-level fundamental laws of nature. Virtually all other kinds of supervenience relations, he argues, are explainable rather than being nomologically basic. But in the case of physical/phenomenal supervenience, he maintains, the reason we cannot imagine what such explanations would be like is that there are no such explanations to be had.

<sup>8</sup>Arguably, even empirical necessary truths like "Water is H<sub>2</sub>O" are so grounded. It is because our concept of water is a *natural-kind* concept, and because natural-kind concepts have a certain distinctive sort of quasi-indexical linkage to the users' environment, that our term 'water' designates H<sub>2</sub>O in all possible worlds.

theoretically costly, and *prima facie* extremely implausible. The Problem of the Many is an unimaginability phenomenon with enormous philosophical weight. Adults should take it very seriously, as grounds for skepticism about the epistemic view of vagueness.

On the other hand, cost–benefit analysis with respect to competing philosophical theories is a complex and multi–faceted business. This is so for vagueness at least as much as for other topics that exercise philosophers. Although the theoretical costs associated with epistemicism are very high indeed, we should bear in mind two kinds of additional factors also at work in philosophical debates about vagueness: the substantial theoretical benefits of the epistemic account of vagueness (e.g., its retention of classical logic and semantics), and the substantial theoretical costs associated with various competing accounts (e.g., the difficulties such accounts typically face in trying to accommodate higher–order vagueness, and the closely related difficulties of avoiding one or another kind of implausible precisification).<sup>9</sup> Williamson’s work on vagueness rightly emphasizes the extent to which these two factors tend to tip the scales back toward epistemicism. As adults we certainly should acknowledge the theoretical virtues of epistemicism, even those of us who find ourselves incapable of believing it.<sup>10</sup>

---

<sup>9</sup>Artificial precisification, with the attendant failure to accommodate higher–order vagueness, arises in standard degrees–of–truth approaches: the successive statements in a sorities series are all assigned precise degrees of truth, which results in a precise boundary between full–fledged truth and the lack thereof and another precise boundary between full–fledged falsity and the lack thereof. Likewise for standard supervaluationist approaches, which are committed to precise boundaries, in a sorites series, between truth *simpliciter* (i.e., truth in all admissible precisifications) and the lack thereof, and between falsity *simpliciter* and the lack thereof. I myself advocate a position that I think solves these problems; see my “Robust Vagueness and the Forced–March Sorites Paradox”, *Philosophical Perspectives* 8, 1994, 159–88; and “Transvaluationism: A Dionysian Approach to Vagueness,” *Southern Journal of Philosophy* 33, 1995, Spindel Conference Supplement, 97–125. My approach incurs high theoretical costs of other kinds, however. Any account of vagueness is likely to be unintuitive and theoretically expensive in one way or another.

<sup>10</sup>My thanks to John Tienson, Mark Timmons, and Tim Williamson for comments and discussion.