

TRANSMUTATIONISM ABOUT VAGUENESS: A PROGRESS REPORT

TERRY HORGAN

ABSTRACT: The philosophical account of vagueness I call “transvaluationism” makes three fundamental claims. First, vagueness is logically incoherent in a certain way: it essentially involves mutually unsatisfiable requirements that govern vague language, vague thought-content, and putative vague objects and properties. Second, vagueness in language and thought (i.e., semantic vagueness) is a genuine phenomenon despite possessing this form of incoherence—and is viable, legitimate, and indeed indispensable. Third, vagueness as a feature of objects, properties, or relations (i.e., ontological vagueness) is impossible, because of the mutually unsatisfiable conditions that such putative items would have to meet. In this paper I set forth the core claims of transvaluationism in a way that acknowledges and explicitly addresses a challenging critique by Timothy Williamson of my prior attempts to articulate and defend this approach to vagueness. I sketch my favored approach to truth and ontological commitment, and I explain how it accommodates the impossibility of ontological vagueness. I argue that any approach to the logic and semantics of vagueness that both (i) eschews epistemicism and (ii) thoroughly avoids positing any arbitrary sharp boundaries (either first-order or higher-order) will have to be not an alternative to transvaluationism but an implementation of it. I sketch my reasons for repudiating epistemicism. I briefly describe my current thinking about how to accommodate intentional mental properties with vague content within an ontology that eschews ontological vagueness. And I revisit the idea, which played a key role in my earlier articulations of transvaluationism, that moral conflicts provide an illuminating model for understanding vagueness.

Terry Horgan is Professor of Philosophy at the University of Arizona and was a member of the University of Memphis Philosophy Department for twenty-one years. He is author (with John Tienson) of *Connectionism and the Philosophy of Psychology* (MIT, 1996) and (with Matjaž Potrč) of *Austere Realism: Contextual Semantics Meets Minimal Ontology* (MIT, 2008). He has published (often collaboratively) many papers on various aspects of philosophy—including philosophy of mind, metaphysics, philosophy of language, epistemology, and metaethics. He edited or coedited several Spindel Conference Supplements to *The Southern Journal of Philosophy*, and contributed the following papers to SJP Supplements: “Supervenience and Cosmic Hermeneutics” (1984), “Settling into a New Paradigm” (with John Tienson, 1988), “Transvaluationism: A Dionysian Approach to Vagueness” (1995), and “What is A Priori and What Is It Good For?” (with David Henderson, 2000).

The philosophical account of vagueness I call “transvaluationism” makes three fundamental claims. First, vagueness is logically incoherent in a certain way: it essentially involves mutually unsatisfiable requirements that govern vague language, vague thought-content, and putative vague objects and properties. Second, vagueness in language and thought (i.e., semantic vagueness) is a genuine phenomenon despite possessing this form of incoherence—and is viable, legitimate, and indeed indispensable. Third, vagueness as a feature of objects, properties, or relations (i.e., ontological vagueness) is impossible, because of the mutually unsatisfiable conditions that such putative items would have to meet.

Williamson (1994) applies the label ‘nihilism’ to philosophical views—my own included—that attribute logical incoherence to vagueness. Being called a nihilist was what prompted me to name my view ‘transvaluationism’. Nietzsche famously held that we could overcome nihilism by embracing what he called “the transvaluation of all values;” similarly, I maintain that vagueness is a genuine and legitimate phenomenon despite harboring a specific kind of logical incoherence. The name ‘transvaluationism’ is meant to evoke a Nietzschean nihilism-transcending attitude toward this inherent logical incoherence—and also seems an appealing label in light of the fact that one popular approach to vagueness is supervaluationism.

In this paper I provide a unified overview of my current thinking about certain philosophical issues involving vagueness, with attention to some salient differences from my early articulation of transvaluationism (see Horgan 1995).¹ In section 1, I set forth the core claims of transvaluationism in a way that acknowledges and explicitly addresses a challenging critique by Williamson (2002) of my prior attempts to articulate and defend this approach to vagueness. (His fundamental charge is that on one natural construal, my account would render impossible not only ontological vagueness but semantic vagueness as well; that is, it would render vagueness impossible *tout court*.) In section 2, I sketch my favored approach to truth and ontological commitment, and I explain how it accommodates the impossibility of ontological vagueness. (Although truth is correspondence to mind-independent reality, very often it is an *indirect* form of correspondence that does not incur ontological commitment to vague objects or vague properties.) In section 3, I argue that any approach to the logic and semantics of vagueness that both (i) eschews epistemicism and (ii) thoroughly avoids positing any arbitrary sharp boundaries (either first-order or higher-order) will have to be, not an alternative to transvaluationism, but an implementation of it. In

¹ I draw in part on material from previous papers, reorganized, modified, and expanded in various ways. See especially Horgan 1997, 2006; material from these two papers is reprinted here with permission of the editors.

section 4, I sketch my reasons for thinking that epistemicism, despite its many theoretical virtues, is surely false. In section 5, I describe a problem that I think arises for any position (including mine) that denies ontological vagueness—namely, the problem of accommodating intentional mental properties with vague content within an ontology that eschews ontological vagueness—and I briefly describe my current ideas about how to address this issue. Finally, in section 6, I revisit an idea that figured centrally in Horgan 1995 but has not been deployed in my more recent writings on vagueness—namely, moral conflicts as a model for the nature of vagueness. I explain, in light of my current articulation of transvaluationism, why and how I think this idea is still useful and illuminating, especially with respect to the question of how a cooperatively motivated subject should behave when being subjected to what I call a “forced march” through a sorites sequence.

1. ARTICULATING TRANSVALUATIONISM

My earliest presentations of the general position I espouse (Horgan 1990, 1994b) predate my use of the label ‘transvaluationism’, which was employed first in Horgan 1995 and thereafter in Horgan 1998, 2000, and 2006. Here I present the core ideas in a way that incorporates and addresses the critique of my views by Williamson (2002). His discussion reveals that my prior attempts to articulate my position were obscure in key respects and also could be construed as embodying certain confusions or conflations—for example, a putative conflation between vague thought and vague discourse being *governed* by mutually unsatisfiable normative requirements on one hand, and their actually *satisfying* such requirements on the other hand. The critique also underscores the looming threat that a version of the view that avoids obscurities and conflations will render any kind of vagueness impossible—not just ontological vagueness, but also vagueness in language and thought.

1.1 *Vagueness as Boundarylessness*

I begin by discussing features that I maintain are essential aspects of any form of vagueness, either semantic or (if such there be) ontological. One is *sorites susceptibility*: for any vague constituent of language or thought, and for any putative vague object, property, or relation, there is a *sorites sequence* directly involving that item. A sorites sequence is a progression of statements, or of states of affairs (actual or possible), that generate a sorites-paradoxical argument. For example, here is a sorites sequence for the vague predicate ‘bald’

and for the putative vague property *baldness* (with ‘ $B(n)$ ’ symbolizing ‘A man with n hairs on his head is bald’):

$$B(0), B(1), \dots, B(10^7)$$

Here is a corresponding sorites-paradoxical argument:

$$(n)[B(n) \supset B(n+1)]$$

$$B(0)$$

$$\text{Therefore, } B(10^7)$$

The claim is that there will always be such a sorites sequence, for any vague linguistic expression, any vague concept, or any putative vague object, property, or relation.

A second essential feature of vagueness is what I call “boundarylessness”—a term I adopt from Sainsbury (1990). This feature, which obtains with respect to a sorites sequence, involves the simultaneous satisfaction by the sequence of the following two conditions:

The Difference Condition: Initially in the sorites sequence there are items with a specific status, and every predecessor of an item with this status has the same status. Eventually in the sequence there are items with the polar-opposite status, and every successor of an item with this status has the same status. No item in the sequence has both the initial status and the polar-opposite status.

The Transition Condition: There is no determinate fact of the matter about status-transitions in the sorites sequence.²

Examples of polar-opposite statuses are baldness vs. nonbaldness, heaphood vs. nonheaphood, satisfying the predicate ‘is bald’ vs. satisfying the expression ‘is not bald’, truth vs. falsity.

The Transition Condition needs further conceptual unpacking. It involves, essentially, two conceptual aspects or dimensions, one individualistic and the other collectivistic.

² In Horgan 1998 I characterized boundarylessness in the following way, concerning sorites sequences of *statements*: “To say that vagueness is boundarylessness is to say that in such a sequence, (1) initially there are true statements (with each predecessor of any true statement being true), (2) later there are false statements (with each successor of a false statement being false), and (3) there is no determinate fact of the matter about the transition from true statements to false statements” (314). My formulation in the current paper is more general, rather than applying merely to statements and their truth values. Williamson (2002) focuses on the earlier formulation, although I think his discussion is applicable, *mutatis mutandis*, to the current formulation too.

The Individualistic Same-Status Principle (ISS Principle): Each item in the sorites sequence has the same status as its immediate neighbors.

The Collectivistic Status-Indeterminacy Principle (CSI Principle): There is no correct overall distribution of statuses to the items in the sequence.

The ISS Principle is so called because it involves items in the sequence considered individually—each considered in relation to its immediate neighbors. The CSS Principle is so called because it involves the items in the sequence considered collectively. Both principles are essentially involved in the idea of boundarylessness—the idea of an absence of sharp boundaries.

Suppose, for instance—contrary to the ISS Principle—that there is some item in a sorites sequence that does not possess the same status as one of its immediate neighbors. Then there is a sharp status-transition between these two items, a sharp boundary—which goes contrary to the very idea of boundarylessness (i.e., the absence of sharp boundaries). The status-transition might not be from the initial status to its polar-opposite status, of course, because there might be one or more intermediate statuses—for instance, an in-between status one might express as “indeterminate whether bald or not bald.” But the point is that a sharp status-transition of *any* kind would violate boundarylessness, even if the transition involves some intermediate status between the original status and its polar opposite.

Suppose now—contrary to the CSI Principle—that there is some correct overall distribution of statuses to all the items in the sorites sequence. Then, given that any sorites sequence must satisfy the Difference Condition, there are bound to be sharp status-transitions between certain successive items in the sequence—because eventually there are items that have the polar-opposite status of the initial items, and no item can have both the initial status and its polar-opposite status. But again, sharp status-transitions of any kind would violate boundarylessness, which is essential to vagueness. So the ISS Principle and the CSI Principle are both aspects of boundarylessness: they both figure as conceptual dimensions in the notion “no fact of the matter about status-transitions.” And boundarylessness is itself the very essence of vagueness.

1.2 *Boundarylessness as Weakly Logically Incoherent*

Since vagueness essentially involves boundarylessness, and the ISS Principle and the CSI Principle are both aspects of boundarylessness, vagueness thereby exhibits a certain specific kind of logical incoherence: namely, the *mutual unsatisfiability* of various principles that are all essential to vagueness. For instance, the Difference Condition and the ISS Principle cannot be jointly satisfied, because the only way for the ISS Principle to be satisfied by a sorites

sequence would be for all items in the sequence to have the same status—contrary to the Difference Condition. Also, the ISS Principle and the CSI Principle cannot be jointly satisfied, given that some items in the sequence have a status (as required by the Difference Condition); for satisfaction of the ISS Principle would require statuses (indeed, the same status) for all the items in the sequence—contrary to the CSI Principle.

But although the mutual unsatisfiability of operative status principles is indeed a form of logical incoherence, it does not necessarily bring in its wake the kind of logical incoherence that generates rampant commitment to contradictions.³ The latter, which I call “strong” logical incoherence, is anathema: language and thought nihilistically deconstruct themselves, once they become enmeshed in rampant contradictions. But the former kind of logical incoherence, I maintain, is “weak”—in the sense that its presence in vague thought and discourse *does not* yield any commitment to rampant contradictions. (Weak logical incoherence does render ontological vagueness impossible, however—a theme to which I return below.) Rather, in semantically proper deployment of vague concepts and vague language, the logical incoherence endemic to vagueness gets quarantined—logically disciplined, held in check—in a manner that effectively fends off strong logical incoherence.

How should one account for this logical discipline? Various options arise here, and the generic transvaluationist framework is neutral among them. From the transvaluationist perspective, many theoretical approaches that have been proposed in the philosophical literature on vagueness—in particular, many approaches involving proposed deviations from classical two-valued logic and semantics—amount to alternative potential implementations of the transvaluationist picture. That is, each such proposal provides some way of quarantining the weak logical incoherence of vagueness, thereby preventing it from generating a logical commitment to rampant contradictions. How a given proposal accomplishes such quarantining will be closely connected to how it handles paradoxical sorites arguments—which is not surprising, since the unchecked use of sorites arguments would be one way of generating contradictions galore. (From the transvaluationist perspective, the sorites paradox brings explicitly into the open the weak logical incoherence endemic to vagueness.)

Consider, for instance, what I call “iterated supervaluationism”—the kind of supervaluationism that treats the metalinguistic category “permissible interpretation” as itself vague, which means that the metalanguage in which

³ The reason I characterize strong logical incoherence in terms of *rampant* commitment to contradictions is that “paraconsistent” systems of logic and semantics allow limited commitment to contradictions. Some philosophers have proposed paraconsistent treatments of vagueness, for instance, Hyde (1997).

supervaluationist semantics is given is subject to a supervaluationist treatment within a second-order metalanguage, and so on up through the hierarchy of metalanguages, *ad infinitum*.⁴ Let ‘ $B(n)$ ’ again symbolize ‘A man with n hairs on his head is bald’, and consider the following sorites-based attempt to derive a contradiction:

$$(1) \quad (n)[B(n) \supset B(n+1)]$$

$$(2) \quad B(0)$$

Therefore,

$$(3) \quad B(10^7) \quad \text{From (1) and (2)}$$

$$(4) \quad \sim B(10^7)$$

Therefore,

$$(5) \quad B(10^7) \& \sim B(10^7) \quad \text{From (3) and (4)}$$

The advocate of iterated supervaluationism will block the sorites subargument for (3), and thereby will block the overall argument for the contradiction (5), by affirming the negation

$$(6) \quad \sim(n)[B(n) \supset B(n+1)]$$

and also affirming the logically equivalent statement

$$(7) \quad (\exists n)[B(n) \& \sim B(n+1)]$$

but nonetheless honoring boundarylessness by refusing to affirm any specific instantiation of (7). And these dialectical moves are entirely appropriate under supervaluationism, because statement (7) is (super)true but no statement of the form $(B(n) \& \sim B(n))$ is (super)true. (On supervaluationist semantics, what would constitute a violation of boundarylessness is not the (super)truth of (7), but rather the (super)truth of one of its instances.)

In Horgan 1994b I sketched an alternative, nonsupervaluationist, approach to the logic and semantics of vagueness that constitutes another potential implementation of transvaluationism. (This approach too must apply iteratively throughout the metalinguistic hierarchy, in order to fully

⁴ The noniterated kind of supervaluationism, in which the metalanguage is governed by classical two-valued logic and semantics, blatantly disrespects the boundarylessness of vagueness: it is committed to sharp boundaries for the metalinguistic category, “permissible interpretation,” and thereby is committed to sharp boundaries between the three categories, “(super)true,” “(super>false,” and “neither (super)true nor (super>false.”

respect the boundarylessness of vagueness.) The basic ideas are the following. First, treat statements like (1), (6), and (7) as neither true nor false, and do the same for their negations. Second, introduce an additional, nonclassical, form of negation, symbolizable by ‘ \neg ’, which works this way semantically: when a statement that is neither true nor false is nonclassically negated, then the resulting negation is true. Thus, these four statements are all true:

$$(8) \quad \neg(n)[\mathbf{B}(n) \supset \mathbf{B}(n+1)]$$

$$(9) \quad \neg(\exists n)[\mathbf{B}(n) \& \sim \mathbf{B}(n+1)]$$

$$(10) \quad \neg \sim(n)[\mathbf{B}(n) \supset \mathbf{B}(n+1)]$$

$$(11) \quad \neg \sim(\exists n)[\mathbf{B}(n) \& \sim \mathbf{B}(n+1)]$$

Third, weak negations operate logically in the following way: if Φ is a logically complex statement, then from the statement $\neg\Phi$ one cannot infer any statement obtained from $\neg\Phi$ by transforming it in a way that “drives the initial negation-symbol inward.” This third feature has the effect of logically “isolating” a statement Φ that is neither true nor false, within the scope of ‘ \neg ’, so that neither Φ nor $\neg\Phi$ can be used inferentially to make trouble. In particular, Φ cannot be used inferentially to generate rampant contradictions; nor can $\neg\Phi$ be so used, because no logically troublemaking statement can be inferred from $\neg\Phi$ by driving the initial negation-symbol inward. The device of nonclassical negation, thus employed, is intended to exert sufficient logical discipline to prevent the weak logical incoherence of vagueness from resulting in strong logical incoherence. The sorites paradox is also handled in an intuitively appealing way: on one hand, statement (1)—which, if accepted as true, would lead by valid reasoning to the false conclusion $\mathbf{B}(n^7)$ —is rejected as neither true nor false; on the other hand, statement (7)—which on this approach counts as asserting the existence of a sharp boundary for baldness (just as it seems to, intuitively)—is also rejected as neither true nor false.

Which of the two lately sketched approaches is better for handling the logic and semantics of vagueness and whether there is yet some further implementation that is better than either of these two are subsidiary issues within the generic transvaluationist perspective—albeit important ones.⁵ Under

⁵ Perhaps there is not, and need not be, any *single* system of logic and semantics that best implements transvaluationism and/or best reconstructs semantically correct use of vague concepts and terms. Also, perhaps transvaluationism can even be implemented by standard two-valued logic, as long as it is employed in a way that respects in practice the logically disciplined weak incoherence of vagueness. Concerning our accommodation of vagueness, Quine (1995) intriguingly remarks, “What I call my desk could be equated indifferently with countless almost coextensive aggregates of molecules, but I refer to it as a unique one of

transvaluationism these alternative approaches remain well worth investigating and well worth subjecting to comparative cost-benefit evaluation, but they all get regarded in the distinctively transvaluationist way: not as avoiding or preventing weak logical incoherence, but rather as exerting logical discipline over it in order to fend off strong logical incoherence.

1.3 *The Impossibility Threat*

If vagueness really does essentially involve mutually unsatisfiable requirements for the items in a sorites sequence, then how could vagueness ever occur at all, even semantically (let alone ontologically)? This question articulates what I will call the *impossibility threat*. Let me spell out the threat in a bit more detail. I claimed above that boundarylessness is essential to vagueness and that boundarylessness is a matter of the Difference Condition and the Transition Condition both being *satisfied* by a sorites sequence. I also claimed that two correlative aspects of the Transition Condition are the ISS Principle and the CSI Principle. Doesn't this mean that in order for the Transition Condition to be satisfied by a sorites sequence, the ISS Principle and the CSI Principle must *themselves* be satisfied by the sequence too? If so, then vagueness is simply impossible across the board—because no sorites sequence can simultaneously satisfy the Difference Condition, the ISS Principle, and the CSI Principle. But if not, then what would it *mean* for the Transition Principle to be satisfied?

The impossibility threat, as just spelled out, is what I take to be the principal thrust of the critique of my prior work on vagueness by Williamson (2002).⁶ Well, I had better not concede that satisfaction of the Transition Condition requires *satisfaction* of both the ISS Principle and the CSI Principle—since that would indeed render impossible any kind of vagueness (given that the Difference Condition must be satisfied too). My task, then, is to explain how satisfaction of the Transition Condition could *involve* both the ISS Principle and the CSI Principle in some essential way—a way other than these latter principles actually being satisfied themselves. I turn to that task in

them, and I do not and cannot care which. Our standard logic takes this . . . in stride, imposing a tacit fiction of unique though unspecifiable reference" (57).

⁶ Williamson (2002) writes: "[T]ransvaluationism avoids inconsistency only if being governed by a semantic standard does not entail satisfying it" (274). Later he says, concerning claims (1)–(3) in my characterization of boundarylessness (cited in note 2 above): "[W]hen Horgan introduces the semantic standards in question, (1)–(3), he treats them as though they were conditions which vague discourse does in fact satisfy, rather than conditions which it somehow ought to satisfy. In particular, he writes as though his point were that there is no determinate fact of the matter about the transition from true statements to false statements, rather than that there ought not to be such a determinate fact of the matter" (283).

the next section, focusing only on vagueness in language and in thought. The topic of ontological vagueness I leave for later in the paper.

1.4 *Vagueness in Language and Thought: Normative Governance by Mutually Unsatisfiable Status Principles*

The key to seeing how the ISS Principle and the CSI Principle could be suitably involved in the Transition Condition's being satisfied in language and thought, without themselves being satisfied, is to bear in mind a distinction. On one hand are principles pertaining to the *statuses* of items in a sorites sequence; these I will call "status principles." (The ISS Principle and the CSI Principle are both status principles.) On the other hand are normative standards for semantically correct judgmental and affirmatory *practice*, in thought and language; these I will call "practice standards." My core claim, couched in terms of this distinction, is as follows. Vagueness in language and thought is possible because the operative semantic practice standards governing vague discourse and vague thought-content are mutually obeyable, in judgmental/affirmatory practice; yet these practice standards, despite being obeyable, nonetheless reflect the normative governance of semantic status principles that are mutually unsatisfiable by the items in a sorites sequence.

What are these mutually obeyable semantic practice standards, and in what sense do they manifest normative governance of judgmental/affirmatory practice by mutually unsatisfiable semantic status principles? Let me address these questions in order. The most pertinent practice standards are a pair of *prohibitions* to which semantically correct judgmental/affirmatory practice must conform—standards that prohibit *assigning* statuses to items in a sorites sequence in a manner that would be in conflict with the ISS and CSI Principles. The prohibitions are these:

The Individualistic Status-Attribution Prohibition (ISA Prohibition): Never attribute a specific status to an item in a sorites sequence and also attribute a different, incompatible, status to its immediate neighbor.

The Collectivistic Status-Attribution Prohibition (CSA Prohibition): Never affirm any determinate overall assignment of statuses to the items in a sorites sequence.

These two practice standards are *mutually obeyable* in semantically correct judgmental/affirmatory practice—and, moreover, are mutually obeyable along with correctly assigning statuses and polar-opposite statuses to *some* items in a sorites sequence.⁷ Thus, semantically correct judgmental/

⁷ Another practice standard that is also in play, and is associated with the status principle I called the Difference Condition, is the following *Predecessor/Successor Requirement*. If one assigns an

affirmatory practice remains possible on the part of agents who think and speak using vague concepts and vague language, even though the ISS and CSI Principles are not mutually satisfiable by the respective items in a sorites sequence.

What constitutes normative governance of judgmental/affirmatory practice by the ISS and CSI Principles? Such governance consists in two correlative facts. First is the fact that these two status principles are *respected in practice* by agents who make judgments and affirmations in accord with the ISA Prohibition and the CSA Prohibition. That is, when an agent's practice conforms to these two practice standards, the agent thereby avoids status assignments that are *in conflict with* the ISS Principle or the CSI Principle.

This feature of competent judgmental/affirmatory practice does not by itself, however, constitute normative governance by the ISS and CSI Principles. For if vague concepts and terms had sharp but unknowable transition-boundaries, then a semantically competent agent would be required to conform in judgmental/affirmatory practice to the ISA and CSA Prohibitions anyway, but for epistemic reasons rather than semantic ones: the agent would never be epistemically warranted in affirming any sharp transitions or any overall status assignment—even though some overall status assignment, containing sharp status-transitions, would be *semantically* correct. (Normatively correct judgmental/affirmatory practice is subject to epistemic standards, not just purely semantic ones.)

So a second fact, also partly constitutive of what it is for the ISS and the CSI Principles to normatively govern proper judgmental/affirmatory practice, is that there *do not exist* any sharp transition-boundaries. (The non-existence of such boundaries is an essential aspect of the notion of boundarylessness—as the term itself indicates.) Thus, the two prohibitory practice standards are grounded not epistemologically but, rather, semantically: not in ignorance of sharp boundaries but, rather, in the ISS and CSI Principles themselves. These status principles are semantically “in force,” even though they cannot be mutually satisfied by the respective items in a sorites sequence. They are in force in the sense that it would be *semantically* incorrect—rather than epistemically incorrect—to engage in judgmental/affirmatory practice that violates them. That is, it would be semantically

initial nonintermediate status to some item in a sorites sequence, one is thereby committed to assignments of that same status to all predecessors of that item in the sequence, and one's judgmental/affirmatory practice must conform to this commitment; likewise, if one assigns the polar-opposite status to some item in the sequence, one is thereby committed to assignments of that same polar-opposite status to all successors of that item, and one's practice must conform to this commitment. This requirement, the ISA Prohibition, and the CSA Prohibition are all three mutually obeyable.

incorrect to assign statuses to items in a sorites sequence in ways that conflict with the ISS Principle or the CSI Principle. Such a status-assignment practice would be semantically incorrect *because* that assignment would be a violation of the ISS Principle and the CSI Principle.⁸

A crucially important feature of both the ISA Prohibition and the CSA Prohibition, of course, is that they can indeed be consistently obeyed in practice—and, moreover, can be consistently obeyed while also assigning statuses to *some* items in a sorites sequence and also assigning polar-opposite statuses to other items. The trick is that in order to exhibit semantically correct judgmental/assertoric practice, one must refrain from undertaking to assign statuses to *all* the items in the sequence—and one must also refrain from ever assigning a specific status to some item and also assigning some different status to an immediate neighbor. (If one is being led on a “forced march” through a sorites sequence [cf. Horgan 1994b], whereby one is asked successively to pronounce on the status of each successive item, one will have to refrain at some point, while simultaneously declaring that one’s choice of such a point does not mark a status-transition.⁹)

Thus, it is indeed possible to engage in affirmatory practice, in thought and language, that is *normatively governed* by the ISS Principle and the CSI Principle—even though it is not possible for these two principles to be simultaneously *satisfied* by one’s thoughts or one’s statements. Semantically correct affirmatory practice involving thoughts that deploy vague concepts, and involving statements that express such thoughts using vague words, can and does exhibit such normative governance; thus, vagueness in language and thought is both possible and actual.

⁸ Are there other plausible cases, apart from vagueness and apart from semantic normativity more generally, in which practice is normatively governed by mutually unsatisfiable requirements or principles? I would say yes, although the cases that come to mind are subject to philosophical dispute. For instance, I would claim that there are genuine *moral conflicts*—situations in which an agent is subject to several mutually unsatisfiable moral obligations, none of which have defeasibility conditions that are met in the circumstances, and all of which therefore remain in force even though one dominates the other. In such cases, remorse is a rationally appropriate attitude concerning one’s failure to meet the dominated obligation; often one also is morally required to do *something* by way of addressing the unmet obligation—such as making reparations, offering an apology, etc. (I return to moral conflicts in section 6 below.) I also would claim that there can be situations in which an agent is subject to several mutually unsatisfiable *epistemic* norms all of which remain in force in the circumstances. For instance, arguably it is rational to accept both quantum mechanics and general relativity, on the grounds of the enormous predictive accuracy and explanatory power of these two theories, even though they are known to be logically incompatible with one another; yet accepting them both goes contrary to the epistemic norm prohibiting the joint acceptance of theories known to be logically incompatible.

⁹ For more on how a maximally cooperative person should handle a forced march, see Horgan (2006) and also section 6 below.

1.5 *The Impossibility of Ontological Vagueness*

The remarks in section 1.4 explain why and how I think the impossibility threat can be overcome, as regards vagueness in language and thought. Ontological vagueness is another matter entirely, however. The weak logical incoherence that accrues to boundarylessness, I claim, renders ontological vagueness outright impossible. Let me elaborate.

I claimed above that boundarylessness is essential to vagueness, that boundarylessness requires the simultaneous satisfaction of the Difference Condition and the Transition Condition, and that satisfaction of the Transition Condition requires the “involvement” of two status principles—the ISS Principle and the CSI Principle. Making sense of such “involvement” is a delicate matter, since the ISS Principle and the CSI Principle are mutually unsatisfiable by the items in a sorites sequence. In the case of semantic vagueness, this can be done—by appeal to mutually obeyable standards for judgmental/affirmatory practice.

But when one asks whether there could be vague objects, or vague properties and relations, the situation is very different. The world itself does not engage in norm-governed practice; thus, it would make no sense to posit practice standards that the world itself supposedly obeys. Indeed, the only clear sense that can be made of the idea of “involvement” of the ISS Principle and the CSI Principle, insofar as putative ontological vagueness is concerned, is to construe involvement as outright *satisfaction*: each item in a sorites sequence would have the same status as its immediate neighbors (e.g., baldness, nonbaldness, or borderline-baldness), and collectively there would have to be no overall distribution of statuses exhibited by the items in a sorites sequence. But the ISS Principle and the CSI Principle *cannot* be mutually satisfied, along with the Difference Condition. Therefore, ontological vagueness is impossible: there cannot be vague objects, vague properties, or vague relations.¹⁰

A way to put the point is this. Although language and thought can be weakly logically incoherent, the world itself cannot be. Weak logical incoherence can arise only with respect to norm-governed *practice*: mutually unsatisfiable status principles exert normative authority by way of mutually obeyable practice standards—that is, in the case of semantic vagueness, standards that

¹⁰ I believe that the present paragraph and the preceding one constitute a satisfactory response to the following remarks in Williamson (2002): “[According to Horgan] vague expressions do not *satisfy* ‘the mutually incompatible semantic requirements that are in force in vague discourse’; they are merely *governed* by them. . . . Horgan provides no argument to show that objects, properties, and relations cannot be governed by mutually incompatible requirements without satisfying them. . . . His argument against ontological vagueness trades on a confusion between satisfying and being governed by” (281–82).

prohibit making status assignments in ways that would violate the status principles. Since the world itself does not engage in norm-governed practice, the world itself cannot exhibit weak logical incoherence. So, the only way for objects, properties, or relations to exhibit boundarylessness would be for them to actually satisfy the ISS Principle and the CSI Principle, along with the Transition Principle. But that is impossible.

It bears emphasizing how far-reaching are the implications of the impossibility of ontological vagueness. The right ontology, whatever it is, cannot include entities answering to thought and talk of “ordinary objects” like tables, chairs, and persons. For, if there were such entities, they would be ontologically vague in certain respects—for example, vague diachronically with respect to temporal boundaries, and vague synchronically with respect to spatial boundaries and physical composition (cf. Van Inwagen 1990). (The same goes for numerous scientific posits both large and small—for example, galaxies, ecosystems, cells, electrons.) But ontological vagueness is impossible, because of the weak logical incoherence of boundarylessness. Hence, the right ontology does not include “ordinary objects,” or persons, or a variety of scientifically posited objects.¹¹

2. DOING WITHOUT ONTOLOGICAL VAGUENESS: TRUTH AS INDIRECT CORRESPONDENCE

Although there is no way around this awkward-seeming metaphysical conclusion, given the boundarylessness of vagueness and the weak logical incoherence of boundarylessness, there is a way of coming to terms with it philosophically. I advocate situating transvaluationism within a wider semantical/ontological framework that I call *contextual semantics*.¹² Very briefly, some leading ideas of contextual semantics are the following. First, truth is *semantic correctness*. (Semantic correctness is distinct from matters of etiquette and is not reducible to epistemic notions such as epistemically warranted assertibility.) Second, very many singular and predicative thought-constituents, together with singular and predicative words of natural language

¹¹ It is sometimes suggested, often with allusion to Lewis (1993), that (i) ordinary objects really do belong to the right ontology, but (ii) they are perfectly precise in all respects, and (iii) there are vastly many more of them than is reflected in ordinary counting practices. For instance, there are vastly many (maybe nondenumerably many) cats on the mat answering to the name ‘Tibbles’, although they largely overlap one another and thus bear the “many but almost one” relation to one another. For an argument that this position is not viable, and also that Lewis (1993) need not be construed as advocating it, see Horgan and Potrč 2008, ch. 5, sec. 5.

¹² See Horgan 1986a, 1986b, 1991, 1995, 2001; Horgan and Timmons 2002; Horgan and Potrč 2006a, 2006b, 2008; and Barnard and Horgan 2006.

that express these thought-constituents, are governed by contextually variable semantic standards involving implicit contextual parameters. (The contextual settings of implicit parameters will normally depend heavily upon one's specific purposes, at the time one deploys the given words or thought-constituents.) Third, in most contexts of language and thought—including numerous contexts of scientific inquiry—the implicit contextual parameters work in such a way that truth is an *indirect* form of correspondence between language and thought on one hand, and the world itself on the other. (The world makes it true, for instance, that Michael Jordan is bald—even though the world itself does not contain any object answering to the name 'Michael Jordan' or any property answering to the predicate 'is bald'.) Fourth, only in certain very unusual contexts of language and thought do the contextually operative semantic standards work in such a way that truth is *direct* correspondence between thought/language and the world. (Contexts of ontological inquiry are an example.)

If contextual semantics is right in its claim that the semantic standards governing thought and language almost always work in such a way that truth is indirect correspondence to the world (rather than direct correspondence), then the numerous vagueness-involving thoughts and statements that we normally take to be true—thoughts and statements deploying singular and predicative constituents that are vague—can really be true after all. They can be true even though the correct ontology does not include any vague objects or vague properties. Given transvaluationism, direct-correspondence semantic standards are very rare indeed. So *of course* it sounds lunatic to say that there are no tables, chairs, or persons—or to say that nothing is a heap and no individual is bald. These claims sound lunatic because in ordinary contexts of usage—contexts governed by indirect-correspondence semantic standards—they would be screamingly false. But in the highly unusual context of ontological inquiry, where one employs direct-correspondence semantic standards in the course of asking about which kinds of objects, properties, and relations are present in the world itself, such bizarre-seeming claims are actually true.

3. WHY NON-EPISTEMICISTS CAN'T HELP BEING TRANSVALUATIONISTS

I explained in section 1.2 why various extant approaches to the logic and semantics of vagueness—especially approaches that both (i) deviate from classical logic and two-valued semantics, and (ii) iterate their nonclassical logico-semantic machinery up through the metalinguistic hierarchy—are

naturally viewed from the transvaluationist perspective not as alternatives to transvaluationism, but rather as different potential implementations of it. In my experience, however, those who pursue these various approaches typically do not acknowledge that vagueness harbors weak logical incoherence; instead they typically suppose (at least implicitly) that their favored approach eschews weak logical incoherence altogether.¹³

Nonetheless, fans of iterated nonclassical approaches usually do have a strong commitment to the following two claims. First, epistemicism is false: vagueness is not a matter of sharp but unknowable boundaries (as epistemicists maintain), but rather is a matter of boundarylessness. Second, the boundarylessness of vagueness not only precludes “first-order” sharp transitions in a sorites sequence (e.g., a sharp transition between truth and falsity, or between baldness and nonbaldness), but also precludes “higher-order” sharp transitions too (e.g., between truth and the status “neither true nor false,” between baldness and the category “borderline bald,” between the category “determinately true” and the category “not determinately true,” etc.). The second claim does indeed reflect the nature of boundarylessness, and this fact underlies the need to iterate one’s nonclassical logico-semantic machinery metalinguistically, rather than articulating that machinery in a metalanguage governed by classical two-valued logic and semantics.

I maintain that anyone who embraces these two claims thereby becomes committed to transvaluationism—whether or not one appreciates this fact or acknowledges it. My argument is abductive, of the form “inference to the only available plausible explanation.” It goes as follows. (For specificity, I formulate it with respect to iterated supervaluationism, but it should be clear enough how it generalizes.) First, a fan of iterated supervaluationism will certainly agree with me that in the case of vagueness, semantically correct

¹³ Brian McLaughlin has suggested to me that the iterated-supervaluationist approach in McGee and McLaughlin 1995 accommodates my own view that there is logical incoherence in vagueness. Several features of their treatment are relevant here. First, they introduce a meta-linguistic ‘definitely’ operator, appendable to the truth predicate. Second, they claim that the pretheoretic notion of truth bifurcates into two distinct notions with respect to vagueness, they express as ‘true’ and ‘definitely true’, respectively. Third, they maintain that the former notion obeys the *disquotation principle*, which “tells us that any adequate understanding of truth ought to give us the [Tarskian] (T)-sentences and (F)-sentences” (214); whereas the latter notion obeys the *correspondence principle*, which “tells us that the truth conditions for a sentence are established by the thoughts and practices of the speakers of a language, and that a sentence is true only if the nonlinguistic facts determine that these conditions are met” (214). McLaughlin’s claim to me is, in effect, that this approach treats the pretheoretic notion of truth as a logically incoherent amalgam of the two distinct truth-notions that he and McGee express as ‘true’ and ‘definitely true’, respectively. But even if they are right that the notion of truth needs to be thus bifurcated, I maintain that the logical incoherence of vagueness cuts deeper. As I am about to argue, logical incoherence is still present within iterated-supervaluationist treatments of vagueness—including theirs, despite its explicit distinction between two notions of truth.

judgmental/affirmatory practice must conform to the two above-described practice standards (the ISA Prohibition and the CSA Prohibition)—and will also agree that the normative grounding for these practice standards is semantic rather than epistemic. (The whole point of *iterating* the supervaluationist machinery, after all, is to avoid sharp semantic boundaries altogether.) But second, fans of iterated supervaluationism have no specific explanation on offer for *why* these practice standards must be adhered to—rather than its being the case that the whole infinite supervaluationist hierarchy of metalanguages fixes a determinate overall assignment of semantic statuses (so that an epistemically ideal agent could then engage in semantically correct practice by actually *assigning* those statuses to the successive statements). *Yet an explanation is needed!* Transvaluationism provides one—namely, that the two prohibitive practice standards (the ISA and CSA Prohibitions) reflect the normative governance of mutually unsatisfiable status principles (viz., the ISS and CSI Principles). Moreover, insofar as one eschews epistemicism and claims that the needed explanation will have to be semantic rather than appealing to putative sharp but unknowable boundaries, this is the *only plausible* (and therefore the best) explanation. So iterated, supervaluationism is really an implementation of transvaluationism rather than a competing position—whether its advocates acknowledge this or not.¹⁴

In order to underscore the need for an explanation of why it is that determinate statuses for items in a sorites sequence do not get fixed by the infinite sequence of truth conditions within the whole iterated supervaluational metalinguistic hierarchy, it will be useful to have an illustration of a case where such fixation *does* occur. For instance, here is an iterated-supervaluationist way of laying down semantics for the predicate ‘Horganic’, as applied to real numbers. The first-order metalanguage M_1 stipulates that in

¹⁴ Williamson (2002, sec. 2) rightly points out that an argument I gave in Horgan 1998, attempting to convict iterated supervaluationists of a commitment to the weak logical incoherence of vagueness, does not work. In essence, I was forgetting that a supervaluationist can consistently—indeed, happily—affirm the following disjunction as (super)true without affirming any of its disjuncts as (super)true:

$$[B(1) \ \& \ \sim B(2)] \vee [B(2) \ \& \ \sim B(3)] \vee \dots \vee [B(10^7 - 1) \ \& \ \sim B(10^7)].$$

I submit that my present argument does the job, and thus constitutes an adequate response to the following remarks from Williamson 2002: “Horgan has therefore failed to show that the iterated supervaluationist is implicitly committed to transvaluationism. Although it would be good news for epistemicism if transvaluationism turned out to be the only alternative, it is too good to be true. Moreover, the failure of Horgan’s attempt also casts doubt on his original argument from boundarylessness to logical incoherence, for iterated supervaluationism provides a way of accepting the premise while denying the conclusion. The appeal to boundarylessness does not motivate the drastic claim of transvaluationism as Horgan had hoped” (279).

any permissible interpretation of the predicate ‘Horganic’, (i) any real number less than 100 belongs to the extension of ‘Horganic’, and (ii) any real number greater than 400 does not belong to this predicate’s extension; nothing else is said in M_1 about this predicate (which effectively allows permissible interpretations to specify any sharp boundary between 100 and 400). The second-order metalanguage M_2 stipulates that in any permissible interpretation I^* of the predicate “permissible interpretation of the predicate ‘Horganic’”, (i) any real number less than 150 belongs to the extension of ‘Horganic’ under any interpretation I of ‘Horganic’ that belongs to I^* , and (ii) any real number greater than 350 does not belong the extension of “Horganic” under any interpretation I of ‘Horganic’ that belongs to I^* ; nothing else is said in M_2 about the first-order metalinguistic predicate “permissible interpretation of the predicate ‘Horganic’” (or about the predicate ‘Horganic’ itself)—which effectively allows permissible interpretations of this higher-order predicate to specify any class of first-order interpretations each of whose member-interpretations specifies a sharp boundary for ‘Horganic’ between 150 and 350. This sequence of stipulations is iterated “Zeno-style” up through the hierarchy of metalanguages *ad infinitum*: the successive lower values increase by increments of half the preceding increment (100, 150, 175, 187.5, . . .), while the successive higher values decrease by the same increments (400, 350, 325, 312.5, . . .). The upshot is that the overall, iterated-supervaluational, semantical hierarchy fixes definite semantic statuses for all statements of the form “Real number r is Horganic.” Such a statement is true for any real number less than 200; is false for any real number greater than 300; and is neither true nor false for any real number that is greater than or equal to 200 and less than or equal to 300.

The moral is that the machinery of iterated supervaluationism does not *by itself* prevent statements in sorites sequences from having determinate semantic statuses; rather, some additional semantic factor is needed. (Accordingly, an expression’s having iterated-supervaluational semantics does not suffice for vagueness; the predicate ‘Horganic’ is not vague.) My claim is that the additional factor is the following, in the case of vagueness: semantically correct judgmental/affirmatory practice is normatively governed by mutually unsatisfiable status principles—namely, the ISS Principle and the CSI Principle. And again, this claim looks to provide the *only plausible* additional factor, insofar as one eschews the epistemicist contention that there are always sharp, albeit unknowable, status-transitions in a sorites sequence.

I can imagine someone responding this way: “The explanation doesn’t require any appeal to weak logical incoherence, but rather is just this: there is *no fact of the matter* about status-transitions in a sorites sequence.” But my reply would be to press hard on the question what is meant here by ‘no

fact of the matter'. It can't just be that the logic and semantics of vagueness is iterated supervenient—as the example of 'Horganic' makes clear. It's boundarylessness, isn't it? But when one focuses carefully on the notion of boundarylessness, doesn't that notion essentially involve the normative governance of both the ISS Principle and the CSI Principle? If not, please tell me *why* not. And please tell me what additional ingredient *does* make for boundarylessness, over and above being subject to iterated-supervenient logic and semantics.

I doubt that satisfying answers to these requests will be forthcoming, or can be. And I suspect that anyone who embraces iterated supervenientism about vagueness, but who denies that vagueness essentially involves weak logical incoherence, simply hasn't thought carefully enough about what could be meant by expressions like 'boundaryless', 'no fact of the matter', and the like as applied to sorites sequences. It's all too easy to be distracted from this crucial question if one is enamored of some fancy formal-semantic machinery as the putative full story about vagueness. Correct implementation of transvaluationism, maybe. Full story, no. The full story, if not epistemicist, must embrace weak logical incoherence.

4. WHY NOT EPISTEMICISM?

Given that a viable non-epistemicist, boundarylessness-respecting account of vagueness evidently needs to be a version of transvaluationism and, therefore, needs to embrace the claim that vagueness essentially involves weak logical incoherence, why not just opt instead for epistemicism? The epistemic theory of vagueness has no need to posit such incoherence, and it also has the significant advantage of retaining classical two-valued logic and semantics.

A reason that might be given for repudiating epistemicism is that if one tries to imagine being able to ascertain the putative sharp boundaries of some vague concept or term, one cannot do so. The thought is this: if epistemicism is true, then even if the capacity to ascertain sharp boundaries would require cognitive abilities beyond those actually possessed by humans, one should at least be able to *imagine* how an idealized cognitive agent without human limitations in discriminatory capacity, memory, etc., would ascertain the putative sharp boundaries of vague categories. Since one cannot imagine this, epistemicism is mistaken.

But Williamson (1997) offers a fairly plausible response to this argument. He assumes, plausibly, that the semantics of vague terms and concepts supervenes on patterns of use—an assumption I propose to grant, for present

purposes.¹⁵ He claims, correctly, that humans cannot ascertain and cognitively survey the overall use-patterns on which the satisfaction conditions for human concepts supervene. Because of this, he claims, even if vague concepts have sharp boundaries, nonetheless humans are not able to know where these boundaries are—and, moreover, are not able even to *imagine* recognizing a boundary transition as a transition. (Humans cannot imaginatively project into the mind of a hypothetical, idealized, cognitive agent who is capable of cognitively surveying the overall use-pattern on which the semantics of terms and concepts supervenes. The requisite projective leap is too great.) Thus, the epistemic theory's account of how vagueness works not only predicts the unknowability of the sharp boundaries possessed by vague categories, but it also predicts that humans cannot even imagine ascertaining these boundaries. Since this fact about unimaginability is predicted by epistemicism, Williamson argues, it does not constitute good evidence against epistemicism.

Let us grant all this, at least for argument's sake (although I myself do find it plausible). For me, epistemicism *still* seems utterly incredible. I think 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. Let me explain.

Consider a snow mass 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 snow mass, 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 snow mass, 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 laws of physics.¹⁶ (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

¹⁵ Someone suspicious of this assumption could replace talk of use-patterns by talk of whatever other kinds of factors one thinks jointly constitute a supervenience base for the semantics of terms and concepts.

¹⁶ Even if, *per impossibile*, we could cognitively survey a complete description of the total intermolecular state of a snow mass, 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.

mystery about avalanches, because we do know the fundamental principles that provide the *explanatory basis* for avalanche phenomena, namely, the relevant laws of physics.

The case of the snow mass 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, delimits 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 “sharp supervenience fact”) that such-and-such specific sharp boundaries supervene on so-and-so overall pattern of usage. In the case of an avalanche, we know what constitutes an explanatory basis of avalanche-onset, namely, the relevant laws of physics, as they apply to the intermolecular state of the snow just before the avalanche. 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-à-vis* putative sharp 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 candidates for sharp boundaries, and there would be no reason whatsoever why one of these candidates should win out over the others.

This deeper problem is an instance of the class of problems that fall under the rubric, introduced by Peter Unger (1980), called “the problem of the many.” Versions of this generic problem exhibit a common form, namely, numerous equally eligible sharp candidates for identity with some item (an object, a meaning, a concept, etc.). The specific version I am stressing here, which I will call “the epistemicist problem of the many” 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 unimaginability featured here 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 snow mass, 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 of the snow mass and the subsequent avalanche, namely, certain physical laws together with those initial conditions.)

The epistemicist problem of the many, I would maintain, is the unimaginability 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 sharp supervenience facts involving vague terms and concepts.
- (2) So there is no such explanatory basis.
- (3) So there are no sharp boundaries for vague terms and concepts.

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

The initial premise cannot be credibly denied, I take it. Moreover, I think the theoretical costs of rejecting the inference from (1) to (2) are unacceptably high, and likewise for the inference from (2) to (3). One way to resist the move from (1) to (2), for instance, would be to claim that sharp supervenience facts have an explanatory basis involving principles that humans are inherently incapable of grasping. Such a view would be the analogue, for vagueness, of Colin McGinn's position concerning the supervenience connections between neurophysical properties and phenomenal mental properties (McGinn 1991, chs. 1–4). 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. As I will put it, he holds that humans are in a state of *deep ignorance* about the explanation.

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 pretheoretic phenomena that need accommodating within an overall philosophical position about mind and brain. The pretheoretic salience of these subjective, experiential, aspects of phenomenal 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 neurophysical 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 pretheoretic phenomena, the situation is quite the reverse: commonsense intuition strongly suggests that vague terms and concepts do *not* have sharp boundaries. Thus, not only is the idea of deep ignorance just as radical *vis-à-vis* vagueness as it is *vis-à-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 wildly implausible.

Suppose an epistemicist questions the inference from (2) to (3). Here the claim would be that although there are sharp supervenience facts, that is, it is indeed the case that sharp 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.”

This too is a very radical position. It entails, for instance, not only that the word ‘heap’ has the same perfectly sharp boundaries in all possible worlds in which ‘heap’ has the same use-pattern as in the actual world (or in all nomologically possible worlds with that use-pattern, if the brute supervenience is nomological rather than metaphysical), but also that *there is no reason* why it is the case that in all these worlds ‘heap’ has just these specific sharp 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 metaphysically necessary truths are grounded in nontendentious *conceptual* truths.¹⁷ 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 sharp boundaries.

¹⁷ Arguably, even empirical necessary truths like “Water is H₂O” 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 H₂O in all possible worlds.

It would be hard enough to believe in such surd modal facts even if sharpness of semantic/conceptual boundaries were part of the pretheoretic phenomena of vagueness. Given that the pretheoretic phenomena instead include the apparent absence of sharp boundaries, the hypothesis of brute supervenience for sharp boundaries seems quite fantastic.

In summary, there is an unimaginability phenomenon associated with vagueness that is deeper than our inability to imagine recognizing a sharp semantic/conceptual transition-point in a sorites series, namely, our inability to imagine any explanatory basis for sharp supervenience facts. This deeper phenomenon yields a powerful-looking argument for the conclusion that vague terms and concepts do not have sharp 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 sharp semantic/conceptual boundaries—is radical, theoretically costly, and *prima facie* extremely implausible. The epistemicist problem of the many provides a powerful reason for rejecting the epistemic view of vagueness.

5. VAGUE INTENTIONAL MENTAL PROPERTIES IN A NON-VAGUE WORLD

I maintain that there is vagueness in thought and language, even though there is no ontological vagueness. I also maintain that the intentionality of thought—its representational content—is more basic than the intentionality of language. Furthermore, I hold that an adequate ontology, although it should eschew any vague objects or properties, nevertheless should recognize and incorporate *mental* properties—including mental properties with vague intentional content. But a *prima facie* worry now arises, which can be put in the following way: Doesn't vague intentional content, as an aspect of a mental property, render that property *ontologically* vague? How then could there be vagueness in thought if indeed there is no ontological vagueness?

One might try finessing this worry by claiming that intentional mental properties should not be, and need not be, countenanced by an adequate ontology. I once proposed such a move myself, under the banner of what I called “preservative irrealism” about mental intentionality (Horgan 1994a). But, under pressure from John Tienson (2002) to provide an ontological foundation for the claim that truth is indirect correspondence, I came to think that the required foundation would need to incorporate mental properties after all—properties that are instantiated either by ontologically precise spatiotemporal locations, or by the whole cosmos in ontologically precise local

manners of instantiation (Horgan and Potrč 2002; 2008, ch. 7). So my present view, in collaboration with Matjaž Potrč, is that the right ontology should include mental properties with vague intentional content.

In order to be consistent with the repudiation of ontological vagueness, this view needs to be understood as claiming that the mental properties in question are *ontologically* non-vague despite being *semantically* vague. That is exactly what I want to claim. But the task is to make sense of such a position. Potrč and I are currently attempting to do so. Our leading ideas are as follows. First, the most fundamental kind of mental intentionality is phenomenologically constituted; it is what lately has come to be called *phenomenal* intentionality. Second, real phenomenal mental properties—the ones that figure in the right ontology—are ontologically precise (non-vague), by virtue of having a precise phenomenal character, a precise overall “what-it’s-like-ness.” (They are non-vague phenomenal “determinate” properties. Although they can fall under vague “determinable” categories, such determinables are not real properties.) Third, having specific vague intentional content can be a phenomenally, and hence ontologically, non-vague aspect of a phenomenal mental property—much as, for example, there can be a phenomenologically precise fact of the matter about the specific extent to which any given region of one’s visual field is *out of focus* at any given moment. The analogy with phenomenally precise degrees of visual blurriness is, as yet, still just an analogy. But it is suggestive, nonetheless. The idea is to situate intentional-mental vagueness within a non-vague ontology, by treating it as an ontologically precise aspect of overall phenomenal character. Vague content can be a phenomenally (and hence ontologically) precise phenomenon, as is fuzzy visual experience.

6. MORAL CONFLICTS AS AN ILLUMINATING MODEL?

In Horgan 1995 I suggested that the workings of vagueness are structurally analogous to situations that are discussed in moral philosophy under the rubric ‘moral conflicts’: the agent is subject simultaneously to two different moral obligations that cannot both be satisfied; both obligations remain in force rather than being defeated (even though one or both of them might have defeasibility conditions that are unmet in the circumstances); yet one of the obligations dominates the other, so that there is an all-things-considered action that is morally best.¹⁸ I said that vagueness involves two conceptual

¹⁸ In Horgan 1995 I described such situations as a species of moral dilemma. But Mark Timmons tells me that what are normally called moral dilemmas in moral philosophy are situations in which two conflicting obligations are both present and neither dominates the other. The situations I had in mind are what are commonly called moral conflicts.

poles, one collectivistic (i.e., pertaining to a sorites sequence as a whole) and the other individualistic (i.e., pertaining to each individual item in a sorites sequence, *vis-à-vis* its immediate neighbors). I said that these two poles involve requirements on sorites sequences that cannot be simultaneously satisfied, and that the collectivistic requirements dominate the individualistic ones without defeating them.

That way of articulating transvaluationism has not been used in my recent writings, particularly since I began to deploy the distinction between status principles on one hand, and practice standards on the other. The question arises, though, whether there might still be a philosophically illuminating analogy between the workings of vagueness and the structure of moral conflicts. It seems to me that the answer is yes. Let me close by briefly elaborating the point.

I have claimed that the two mutually unsatisfiable semantic status principles—the CSI Principle and the ISS Principle—both exert normative governance over thought and discourse employing vague concepts and words, and that such governance is manifested in mutually obeyable, prohibitory, semantic practice standards—the CSA Prohibition and the ISA Prohibition. Is there an interesting sense in which the CSI Principle *dominates* the ISS Principle? Indeed so. Such dominance reveals itself when one considers how someone who is fully cooperatively motivated, and who is fully competent conceptually and semantically, should form successive judgments involving some vague category (say, “bald”) when being force-marched (as I put it) through a pertinent sorites sequence (e.g., a sequence of men each of whom has a tiny bit more hair on his head than his predecessor) and is being asked to form a sincere judgment at each stage about the applicability of the vague category.

In such a forced-march situation, it is possible to correctly assign statuses to items in the sequence in a manner that fully conforms to the CSI Principle: one can do so by assigning statuses to only *some* of the items rather than to all of them, and doing so without ever assigning some status to an item and a different status to one of its immediate neighbors. One *goes silent* on some of the items, without assigning any status to them at all. (Perhaps one even says explicitly, when going silent somewhere in the march or undertaking new classificatory assertions later in the march, that one’s going silent or resuming classificatory responses is not to be construed as indicating that one is marking a semantic transition.) By contrast, it is not possible, in a forced-march situation, to correctly assign statuses in a manner that fully conforms to the ISS Principle; for, doing that would require assigning the same status to every single item in the sorites sequence. Herein we have an important sense in which the CSI Principle dominates the ISS Principle: one can actually *assign* statuses, correctly, in a

manner that conforms to the CSI Principle (although of course one does not assign statuses to *all* the items in the sorites sequence), whereas one cannot correctly assign statuses in a manner that conforms to the ISS Principle (i.e., one cannot correctly assign the same status to each item's immediate neighbors that one assigns to the item itself). In the case of the ISS Principle, rather, the best one could do would be to respect it, by obeying the practice standard I have called the ISA Prohibition. So the analogy with moral conflicts still looks apt and illuminating, even though much of the theoretical weight I earlier put upon this analogy is now borne instead by the distinction between status principles and practice standards.¹⁹

REFERENCES

- Barnard, R., and T. Horgan. 2006. Truth as mediated correspondence. *Monist* 89: 28–49.
- Horgan, T. 1986a. Psychologism, semantics and ontology. *Noûs* 20: 21–31.
- . 1986b. Truth and ontology. *Philosophical Papers* 15: 1–21.
- . 1990. Psychologistic semantics, robust vagueness, and the philosophy of language. In *Meanings and prototypes: Studies in linguistic categorization*, ed. S. L. Tsohatzidis. London: Routledge.
- . 1991. Metaphysical realism and psychologistic semantics, *Erkenntnis* 34: 297–322.
- . 1994a. Naturalism and intentionality. *Philosophical Studies* 76: 301–26.
- . 1994b. Robust vagueness and the forced-march sorites paradox. *Philosophical Perspectives* 8: 159–88.
- . 1995. Transvaluationism: A Dionysian approach to vagueness. *Southern Journal of Philosophy* (Spindel Conference Supplement on Vagueness) 33: 97–126.
- . 1997. Deep ignorance, brute supervenience and the problem of the many. *Philosophical Issues* 8: 229–36.
- . 1998. The transvaluationist conception of vagueness. *Monist* 81: 316–33.
- . 2000. Facing up to the sorites paradox. In *Proceedings of the 20th World Congress of Philosophy. Volume 6: Analytic Philosophy and Logic*, ed. A. Anamori, 99–111. Philosophy Documentation Center.
- . 2001. Contextual semantics and metaphysical realism: Truth as indirect correspondence. In *The Nature of Truth: Classic and Contemporary Perspectives*, ed. M. Lynch, 67–95. Cambridge, MA: MIT Press.
- . 2006. Transvaluationism: The benign logical incoherence of vagueness. *Harvard Review of Philosophy* 14: 20–35.

¹⁹ Thanks to Matjaž Potrč, Mark Timmons, and an audience at the University of Nebraska for helpful discussion, and to Timothy Williamson for being such a helpful philosophical critic. With fond memories of my twenty-one years in the University of Memphis Philosophy Department, I dedicate this paper to Chris Spindel and to the memory of her late husband Murray Spindel.

- Horgan, T., and M. Timmons. 2002. Conceptual relativity and metaphysical realism. *Philosophical Issues* 12: 74–96.
- Horgan, T., and M. Potrč. 2002. Addressing questions for blobjectivism. *Facta Philosophica* 4: 311–21.
- . 2006a. Abundant truth in an austere world. In *Truth and realism: New essays*, ed. P. Greenough and M. Lynch, 137–67. Oxford: Oxford University Press.
- . 2006b. Particularist semantic normativity. *Acta Analytica* 21: 45–61.
- . 2008. *Austere realism*. Cambridge, MA: MIT Press.
- Hyde, D. 1997. From heaps of gaps to heaps of gluts. *Mind* 106: 440–60.
- Lewis, D. 1993. Many but almost one. In *Ontology, causality, and mind: Essays on the philosophy of D. M. Armstrong*, ed. K. Campbell, J. Bacon, and L. Reinhardt. Cambridge: Cambridge University Press.
- McGee, V., and B. McLaughlin. 1995. Distinctions without a difference. *Southern Journal of Philosophy* 33 (Supplement): 203–51.
- McGinn, C. 1991. *The problem of consciousness*. Oxford: Blackwell.
- Quine, W. V. O. 1995. *From stimulus to science*. Cambridge, MA: Harvard University Press.
- Sainsbury, R. M. 1990. *Concepts without boundaries*. Inaugural Lecture, King's College London. Reprinted in *Vagueness: A reader*, ed. R. Keefe and P. Smith. Cambridge, MA: MIT Press, 1996.
- Tienson, J. 2002. Questions for blobjectivism. *Facta Philosophica* 4: 301–10.
- Unger, P. 1980. The problem of the many. *Midwest Studies in Philosophy* 5: 411–67.
- Van Inwagen, P. 1990. *Material beings*. Ithaca and London: Cornell University Press.
- Williamson, T. 1994. *Vagueness*. London and New York: Routledge.
- . 1997. Imagination, stipulation and vagueness. *Philosophical Issues* 8: 215–28.
- . 2002. Horgan on vagueness. In *Essays on the philosophy of Terence Horgan*, ed. J. Brandl and O. Markič, 273–85. Rodopi. Volume also issued as *Grazer Philosophische Studien* 63 (2002).