

The Many Minds Account of Vagueness¹

There are two natural objections to epistemicist accounts of vagueness that one frequently encounters in conversation (objections that are frequently run together).² One objection is that it is hard to live without an informative answer to the question as to how the non-semantic facts – non-relational and relational -- about a given individual determine the semantic profile of that individual. Let us call this the Bruteness worry. A second objection is that it seems metaphysically arbitrary that just one of the many candidate cut-offs for, say, baldness, should be associated with some ordinary pattern of use of the term ‘bald’. Let us call this the Arbitrariness worry. In this paper we sketch an epistemicist approach to vagueness that provides a distinctive perspective on these foundational issues. On the one hand, it affords epistemicism protection against the Arbitrariness worry, by combining the even-handedness typically associated with supervaluationism with the logic and semantics distinctive of epistemicism.³ And on the other, it renders the Bruteness objection less pressing by assimilating it to a kind of bruteness that many of us have already learned to live with. Our view relies on a plenitudinous metaphysics to which we are antecedently sympathetic. Like any metaphysically laden view, it is inevitably somewhat tendentious. But its unity and elegance make it at least deserving of a place on the philosophical map.

¹ We are very grateful to Timothy Williamson for helpful comments.

² The main target of these objections is Timothy Williamson, *Vagueness*, Routledge, 1994.

³ For a presentation of supervaluationism, see Kit Fine, ‘Vagueness, Truth and Logic,’ *Synthese* 30 (1975), p.265-300.

We begin by assuming a liberal ontology, one according to which any natural object is at every time mereologically coincident with myriad other natural objects, some of which exist for the same temporal duration, some of which don't. A lump of clay that is coincident with a statue is also coincident with very many other objects that may vary from the original pair on a variety of parameters, including persistence conditions and modal profile. Positing this plenitude protects one from accusations of anthropocentrism that would inevitably arise were one to think, for example, that there was merely a statue and a lump and nothing else. And, relatedly, it protects one from puzzles of metaphysical arbitrariness that would naturally arise were one to claim that out of all the conceivable persistence conditions and modal profiles and associated essences, only two are actualized in the relevant vicinity. We should mention that standard presentations of both supervenience and epistemicism proceed within a plenitudinous setting⁴, and so the plenitude-driven account of vagueness that follows ought to be particularly palatable to such theorists.

How do we manage to pick out a statue rather than a lump with our term 'That statue' when both are compresent? This question has an easy answer – the conception associated with that complex demonstrative is such that a lump is a poor reference candidate.⁵ Why is it that the statue but not the lump passes out of existence when both are crushed despite the fact that one is microphysically just like the other up to that time? This question also has an easy answer: it is because they have different essences/natures

⁴ Consider standard epistemicist and supervenience approaches to the 'problem of the many', which both rely on positing many candidates in reality. (See, for example, David Lewis, 'Many But Almost One,' in *Papers in Metaphysics and Epistemology*, Cambridge University Press, 1999.)

⁵ Of course there is a trickier question: Granted you don't pick out lump-like things with 'That statue'. But given plenitude and the vagueness of 'That statue' there are many candidate entities for 'That statue' to refer to. How do you pick out one of them. The account of vagueness we offer below is designed to offer an interesting answer to questions like this.

that determine different modal and persistence profiles for them. Why is it that one particular object has the statue essence and another has the lump essence rather than the other way around? This is a confused question. Asking why a thing has its essence rather than another one is like asking why I have my haccuity rather than someone else's.

All views have foundational facts that are not in turn grounded in further, deeper facts. Still, some will find the postulation of multiple coincident objects deeply unsatisfying: 'How can it be just a brute fact that a given thing has the individual essence that it does, a fact that cannot be grounded in its microphysical profile? After all, a microphysical twin with which it is coincident has a different individual essence.' Such philosophers antecedently expected the modal profile of a thing to be fixed by its microphysical profile. But singular facts about individual essences seem like an excellent stopping point in the order of explanation: the relevant bruteness is there, all right, but it is not untoward. This is the position of those, like ourselves, who are willing to distinguish the statue from the lump even in a case where the two coincide for their lifetimes and thus enjoy the same microphysical history.

Coincident things can vary in a number of ways. One of the dimensions of variation is responsible for the phenomenon of vagueness. Or so we would like to suggest. Let the psycho-semantic profile of an object be a function that describes what the object thinks at the various possible circumstances in which it exists. (We can think of it as a partial function from centred worlds to propositions, which delivers no value for centred worlds at which the object does not exist, and delivers the set of propositions that the object would be thinking were the world and time of the centre actual and current.) Now on the bold hypothesis that coincident objects can have slightly different psycho-

semantic profiles, we obtain a rather elegant picture of how the phenomenon of vagueness arises. As a toy picture, imagine that there are three candidate meanings for the predicate 'bald': bald¹, bald² and bald³. Suppose a person says 'Someone is bald'. In reality, there are many coincident thinking things.⁶ Some of them are thinking the proposition that someone is bald¹, some of them are thinking the proposition that someone is bald², some of them are thinking the proposition that someone is bald³. Enlightened by our paper, each of them realizes that it is their nature to be mereologically coincident with many minds. When each of them says 'I am coincident with many minds', each uses that noise token to express the (true) proposition that she is coincident with many minds, though of course there are many minds that each use that very same noise token to each make a different speech act. Suppose, again enlightened by our paper, the coincident minds at a certain location realize that some of them mean bald¹ by 'bald' some of them mean bald² by 'bald' and so on. They utter 'I mean either bald¹ or bald² or bald³ by 'bald' but I don't know which.' It seems clear that each expresses a true proposition about herself by that utterance. And supposing Fred is bald¹ but not bald² or bald³, they may say 'The sentence 'Fred is bald' is true or false. That turns on whether I mean bald¹ or not. But since I don't know whether 'Fred is bald' expresses the proposition that Fred is bald¹, I cannot know whether 'Fred is bald' is true or not.' Some of the coincident minds that utter this mean bald¹ by 'bald', some don't. But each of them uses that speech to make a perfectly accurate self-description.⁷

⁶ Further complications to the basic picture will no doubt be needed to account for the traditional problem of the many, whereby objects with slightly different boundaries but which largely overlap all seem to compete for personhood. From the current perspective it will be natural to think that for each relevant set of atoms, there are multiple fusions, with varying psycho-semantic profiles.

⁷ If a conversation is going on, it is natural to think that people with coordinate meanings are more suitable conversation partners than people who are not. It is thus natural to think that like-minded thinkers are more

We hope the basic picture should now be clear to the reader. The account of vagueness is broadly epistemicist. The logic and semantics for every thinking thing is ordinary classical logic and semantics. Bivalence and excluded middle are thus happily embraced. Further, the proponent of this view claims, with the orthodox epistemicist, that in a case of vagueness we are unable to know the answer. But this view provides a distinctive answer to the question of what distinguishes ignorance due to vagueness from ordinary ignorance (after all, not all ignorance is rooted in vagueness). Ignorance due to vagueness is rooted in the phenomenon wherein coincident things have slightly different semantic profiles. Note that this metaphysical picture makes good on the thought that vagueness is a kind of in-principle ignorance. Given that any de se thought token is the vehicle for many different de se attitudes by many minds, it seems wholly unclear how, even in principle, one of the many minds could know de se of himself that he was thinking bald¹ by ‘bald’. Of course they could utter ‘I mean bald¹ by ‘bald’ and some of them speak the truth. But in doing so many minds coincident with them would speak falsely. The proximity of error, proximity in a very real sense, would seem to rob each truth-teller’s belief from the status of knowledge.⁸ At no metaphysically possible world is this problem overcome. Ignorance that stems from the kind of psycho-semantic variation we have just described is thus an in-principle kind of ignorance.

Let us review the standard kind of Arbitrariness complaint against epistemicism:
‘There are many candidate meanings, bald¹ bald² and bald³. You tell me that exactly one

suitable referents for ‘You’. When ‘You’ is uttered in a conversation, many minds express different singular thoughts, referring to different conversation partners. (We shall not pursue this issue further here.)

⁸ After all the method on which the judgment is based leads to falsehood in relevantly similar cases. The situation is similar to one in which a pair of individuals is told that one of them is Jack Lyons. If Jack Lyons were to think ‘I am Jack Lyons’, simply on the basis of the information, he would be epistemically relevantly wrong. (The status of the proposition he expresses with regard to metaphysical modality is not relevant.)

of them but only of them is picked out by 'bald'. Suppose bald¹ is the winner. It seems metaphysically arbitrary, almost magical, that bald¹ should win the competition. For there seems to be nothing in nature that would determine that bald¹ should be expressed and yet, say, bald² not expressed.'

It is hard to see how this complaint can be effectively adapted to the current proposal. After all, the current proposal does not propose an asymmetrical treatment of bald¹, bald² and bald³ in connection with actuality. It is indeed correct that if all three are candidates then nothing in nature determines that one should be expressed and yet the other not: after all, for each candidate some of the coincident minds express that meaning by 'bald'. In this way, our epistemicism is utterly even handed in its attitude to the candidates.

Of course one could still press the Bruteness worry: Suppose one coincident mind, Bill, expresses bald¹ and another, Ben, expresses bald². Why should the semantic facts be that way rather than the other way around? But this is no more pressing than in the statue lump case. Consider, for example, a statue-lump scenario where one of an erstwhile coincident pair, Jane, survives crushing but the other, Sally, does not. The question as to why things are not the other way around with respect to psycho-semantic profiles of Bill and Ben is like asking why things are not the other way around with respect to persistence profiles of Jane and Sally. In each case, the nature of the objects in question – their persistence conditions in one case and their psycho-semantic profile in the other case – underwrites the actual and possible facts of relevance.

Granted, the many minds view implies that semantic properties cannot be identified with microphysical properties, since Bill and Ben are microphysical duplicates

and not semantic duplicates. But this kind of anti-reductionism should also have been embraced by anyone committed to statue-lump plenitude. After all, the fact that numerically distinct statues and lumps could come in and out of existence at the same time (or else both be eternal ...) shows that the property of being a statue cannot be identified with a microphysical property, not even a highly extrinsic and disjunctive one.⁹ And once such anti-reductionism is accepted, the question ‘Why does Bill have the semantic profile he does rather than, say, Ben’s?’ will be reckoned as misconceived as asking ‘Why does Sally have the persistence profile she does rather than, say, Jane’s’.¹⁰

What about second order vagueness? So far we have spoken of the candidate meanings of a term. But isn’t there some vagueness in the notion of candidacy? Let us sketch one way of thinking about higher order vagueness within the current framework, again with the help of a toy model. Suppose there are two relations, *thinks*¹ and *thinks*². When ‘Someone is bald’ is uttered, there are many coincident entities, some of which stand in the *thinks*¹ relation to a proposition, others stand in the *thinks*² relation to a proposition (though none stand in both relations). Hence the natures of the coincidents determine some to be *think*¹ers, others to be *think*²ers. Suppose that the *think*¹ relation is slightly more flexible than the *think*² relation. Each of *bald*¹, *bald*² and *bald*³ is thought

⁹ Of course, this kind of anti-reductivism is compatible with certain kind of global microphysical supervenience theses. See, for example, "[In Defense of Global Supervenience](#)", with R. Cranston Paull, *Philosophy and Phenomenological Research* 52, 1992, p 833-854

¹⁰ The version of the many minds view that we have presented – and which we are most attracted to – is one according to which psycho-semantic profile constitutes a fundamental parameter of variation -- along with persistence conditions etc. There is another version of the many minds view according to which the semantic profile of an object *does* supervene on its intrinsic and relational non-semantic profile, so long as the latter profile includes temporal and modal facts. On that view, the difference in the semantic profile of ‘bald’ in the mouth of many coincident minds is grounded in slightly different modal/temporal profiles of the many minds. This makes the bruteness worry more pressing: how it is that such subtle variations in dispositional and temporal profiles make the difference than they do? We are not drawn to this view, but we mention it for the sake of comprehensiveness.

about¹ by some or other coincident entity. But only bald¹ and bald² are thought about² at that spatiotemporal location at that world.

Suppose this is how things are. Here is an enlightened speech:

‘I am one of many coincident things. If I am a think¹er, then thinkers are think¹ers. After all if I am a think¹er then I think¹ about think¹ing when I utter ‘think’, and I don’t think² about anything. (I don’t think about the gerrymandered disjunction of think¹ing and think²ing when I think about thinking: natural relations attract but gerrymanded disjunctions repel from a semantic point of view.) If I am a think¹er, then thinking is in one important respect less restrictive than the think² relation in that various entities are thinking about bald³ in this location at this world but none thinks² about bald³. Meanwhile if I am a think²er, then thinking is think²ing, and thinking is an important respect more restrictive than the think¹ relation in that no entity is thinking about bald³ but some entity think¹s about bald³.’

Suppose this speech is produced. Adopting the view from nowhere, what is going on is that many coincident things share in that speech token. Some think¹ a true proposition, others think² a true proposition.

The speech can felicitously continue thus: ‘If I am a think¹er then the candidate meanings for ‘bald’ at this context are the things that are thought by someone at this location at this world – that is all and only bald¹, bald², and bald³. If I am a think²er, then the candidate meanings for ‘bald’ are the things that are thought by someone at this location at this world – that is all and only bald¹ and bald².’ Returning to the view from nowhere: The think¹ers think¹ a true proposition and denote¹ think¹ing by ‘think’. The think²ers think² a true proposition and denote² think²ing by ‘think’. The vagueness in the

range of candidates – and thus the phenomenon of higher order vagueness -- has one again been grounded in subtle variations in the profiles of coincident entities.

Obviously, this metaphysical account of vagueness will, for many, come as a bit of a shock to their sensibilities. But its merits should not be overlooked. In requiring a plenitude of mereologically coincident objects, it merely taps into an altogether respectable tradition. In underwriting classical logic and semantics, it avoids dubious dealings with novel and speculative logico-semantic approaches that we can barely learn to think with and whose very coherence is often open to question. And in its even handedness towards the candidate meanings of 'bald' it avoids the arbitrariness which scares people off the more standards presentations of epistemicism. Moreover, it offers a satisfying account of the in-principle ignorance that many of us feel is a distinctive feature of vagueness, and a way of distinguishing in-principle ignorance due to vagueness from in-principle ignorance that has other sources. Anyone working hard in the field of vagueness ought to expect surprises. For our money, the elegance of our metaphysical proposal makes this surprise a moderately pleasant one.