

Many Many Problems

Recently four different papers have suggested that the supervenient solution to the Problem of the Many is flawed. Stephen Schiffer (1998, 2000a, 2000b) has argued that the theory cannot account for reports of speech involving vague singular terms. Vann McGee and Brian McLaughlin (2000) say that theory cannot, yet, account for vague singular beliefs. Neil McKinnon (2002) has argued that we cannot provide a plausible theory of when precisifications are acceptable, which the supervenient theory needs. And Roy Sorensen (2000) argues that supervenientism is inconsistent with a directly referential theory of names. McGee and McLaughlin see the problem they raise as a cause for further research, but the other authors all take the problems they raise to provide sufficient reasons to jettison supervenientism. I will argue that none of these problems provide such a reason, though the arguments are valuable critiques. In many cases, we must make some adjustments to the supervenient theory to meet the posed challenges. The goal of this paper is to make those adjustments, and meet the challenges.

1. Schiffer's Problem

Stephen Schiffer suggests the following argument refutes supervenientism. The central point is that, allegedly, the supervenient theory of vague singular terms says false things about singular terms in speech reports.

Pointing in a certain direction, Alice says to Bob, 'There is where Harold and I first danced the rumba.' Later that day, while pointing in the same direction, Bob says to Carla, 'There is where Alice said she and Harold first danced the rumba.' Now consider the following argument:

- (1) Bob's utterance was true
- (2) If the supervenient semantics were correct, Bob's utterance wouldn't be true
- (3) ∴ The supervenient semantics isn't correct. (Schiffer 2000b: 321)

Assuming Bob did point in pretty much the same direction as Alice, it seems implausible to deny (1). The argument is valid. So the issue is whether (2) is correct. Schiffer has a quick argument for (2), which I will paraphrase here. On supervenient semantics, a sentence is true iff each of its acceptable precisifications is true. In this case, this means that if Bob's utterance is true then it must be true however we precisify 'there'. Each precisification of 'there' will be a (precise) place, and since 'there' is rather vague, many of these precisifications will be acceptable. For Bob's utterance to be true, then, Alice must have said of every one of those places that it was the place where Harold and her first danced the rumba. But Alice couldn't have said all those things, so (2) is true.

Schiffer suggests that one way out of this problem would be to accept the existence of a vague object: the place where Harold and Alice first danced the rumba. I will note in section four several reasons for thinking the cost of this move is excessive. Fortunately, there is a cheaper way home.

Schiffer underestimates the scope of supervenientism. On Schiffer's vision of the theory, a precisification assigns a precise content to a word, and hence to a sentence, then the world determines whether that content is satisfied, and hence whether the sentence is true on that precisification. This is hardly an unorthodox view of how supervenientism works, it seems for instance to be exactly the view defended in Keefe (2000), but it is neither the only way, nor the best way, forward. We could say, rather, that a precisification assigns content to *every linguistic token in the world*, and the truth conditions of every one of these tokens is then determined relative to that global assignment of content. So if a precisification P assigns a place x to Bob's word 'there', Bob's utterance is true according to that precisification iff P also assigns x to Alice's utterance of 'there'. That is, Bob's utterance is true according to P iff the precisification of his words by P just is what Alice said *according to P*.¹

It is a dramatic widening of the scope of precisifications to claim that they assign content to every linguistic token in the world, rather than just words in the sentence under consideration, but it can be justified.² Consider how we would react if later in the day, pointing in the crucial direction, Alice said, 'Harold and I never danced the rumba there.' We would think that Alice had contradicted herself – that between her two statements she must have said something false. A standard supervenientist account, where sentences are precisified one at a time, cannot deliver this result. On such a view, it might be that each of Alice's utterances are true on some

¹ Following Schiffer, we ignore the vagueness in 'is where Harold and I first danced the rumba.' This phrase is vague, but its vagueness raises no extra issues of philosophical importance.

² Thanks to John Hawthorne for the following argument.

precisifications, so they are both neither true nor false. On my theory, each precisification applies to both of Alice's utterances (as well as every other utterance ever made) and since on each precisification one or other of the utterances is false, it turns out supertrue that Alice said something false, as desired. The current view allows for penumbral connections between sentences, as well as penumbral connections within sentences. Just as someone who says, "That is red and orange" says something false, my theory decrees that someone who says, "That is red. That is orange," while pointing at the same thing says something false, even if the object is in the vague area 'between' red and orange.

It is crucial for this response to work that on every precisification, Alice and Bob's demonstratives are co-referential. It does not seem like a particular expansion of supervenient theory to posit this as a penumbral connection between the two words. At least, it seems plausible enough to do this if Alice and Bob really are pointing in a similar direction. If their demonstrations are only *roughly* co-directional, then on some precisifications they may well pick out different objects. This will definitely happen if some admissible precisification of Alice's 'there' is not an admissible precisification of Bob's 'there'. In such a case, the theory here predicts that Bob's utterance will be indeterminate in truth value. But if Alice and Bob only vaguely pointed in the same direction this is the correct prediction.

2. Natural Properties

Schiffer's problem seems to have been solved with a minimum of fuss, but there is still a little work to do. Above I posited a penumbral connection between Alice's and Bob's words without *explaining* how such a connection could arise. This connection can be explained by some general considerations about content, considerations closely tied to the view of vagueness as semantic indecision that provides the best motivation for supervenientism. As a few writers have pointed out (Quine 1960, Putnam 1980, Kripke 1982), there is not enough in our dispositions to use words to fix a precise content all terms in our lexicon. This does not immediately imply a thorough-going content scepticism because, as a few writers have also pointed out (Putnam 1973, Kripke 1980, Lewis 1983, 1984), meanings ain't (entirely) in the head. Sometimes our words refer to a particular property or object rather than another not because our dispositions make this so, but because of some particular feature of that property or object. David Lewis calls this extra feature 'naturalness': some properties and objects are more natural than others,

and when our verbal dispositions do not discriminate between different possible contents, naturalness steps in to finish the job and the more natural property or object gets to be the content.

Well, that's what happens when things go well. Vagueness happens when things don't go well. Sometimes our verbal dispositions are indiscriminate between several different contents, and no one of these is more natural than all the rest. In these cases there will be many unnatural contents not eliminated by our dispositions that naturalness does manage to eliminate, but there will be still be many contents left uneliminated. Consider, for example, all the possible properties we *might* denote by 'tall woman'. As far as our usage dispositions go, it might denote any one of the following properties: woman taller than 1680mm, woman taller than 1681mm, woman taller than 1680.719mm, *etc.* And it does not seem that any of these properties are more natural than any other. Hence there is no precise fact about what the phrase denotes. Hence it is vague. In sum, our dispositions are never enough to settle the content of a term. In some cases, such as 'water', 'rabbit', 'plus', 'brain' and 'vat', nature is kind enough to, more or less, finish the job. In others it is not, and vagueness is the result.

(The above reasoning has a surprising consequence. Perhaps our verbal dispositions are consistent with the predicate *Tall X* denoting the property of being in the top quartile of *Xs* by height. Unlike each of the properties mentioned in the text, this is a more natural property than many of its competitors. So if this kind of approach to vagueness is right, there *might* not be quite as much vagueness as we expected.)

If this is how vagueness is created, then there is a natural way to understand how precisifications remove vagueness. Vagueness arises because *more natural than* is a partial order on putative contents, and hence there might be no most natural content consistent with our verbal dispositions. If this relation only defined a strict ordering, so whatever the candidate meanings were, one of them would be most natural, vagueness might be defeated. Well, that isn't true in reality, but it is true on each precisification. Every precisification is a completion of the 'naturalness' partial order. That is, each precisification *P* defines a strict order, *more natural-P than*, on possible contents of terms such that α_1 is more natural-*P* than α_2 if (but not only if) α_1 is more natural than α_2 . The particular contents of terms according to *P* is then defined by using the more natural-*P* than relation where the more natural than relation is used in the real theory of content.

This conjecture meshes nicely with my theory of the role of precisifications. First, it explains why precisifications apply to the whole of language. Since a precisification does not just remedy a defect in a particular word, but a defect in the content generation mechanism, precisifications are most naturally applied not just to a

single word, but to every contentful entity. Secondly, it explains why we have the particular penumbral connections we actually have. Recall that it was left a little unexplained above why Alice's and Bob's use of 'there' denoted the same precise place. On the current conjecture, Alice's term refers to a particular place x according to P because x is more natural- P than all the other places to which Alice might have referred. If this is so, then x will be more natural- P than all the other places to which Bob might have referred, so it will also be the referent according to P of Bob's there. Hence according to every precisification, Bob's utterance will be true, as Schiffer required.

We can also explain some other unexplained penumbral connections by appeal to naturalness. Consider the sentence *David Chalmers is conscious*. Unless this is supertrue, supervenience is in trouble. It is vague just which object is denoted by *David Chalmers*. On every precisification, there are other objects that massively overlap David Chalmers. Indeed, these very objects are denoted by 'David Chalmers' on other precisifications. These objects are not conscious, since if one did there would be two conscious objects where, intuitively, there is just one. But each of these rogue objects must be in the extension of 'conscious' on the precisifications where it is the denotation of 'David Chalmers'. So 'conscious' must be vague in slightly unexpected ways, and there must be a penumbral connection between it and 'David Chalmers': on every precisification, whatever object is denoted by that name is in the extension of 'conscious', while no other potential denotata of 'David Chalmers' is in the extension. How is this penumbral connection to be explained? Not by appeal to the meanings of the terms! Even if 'David Chalmers' has descriptive content, it is highly implausible that this includes *being conscious*. (After all, unless medicine improves a bit in a thousand years Chalmers will not be conscious.) Rather, this penumbral connection is explained by the fact that the very same thing, naturalness, is used in resolving the vagueness in the terms 'conscious' and 'David Chalmers'. If the precisification makes one particular possible precisification of 'David Chalmers', say d_1 , more natural than another, d_2 , then it will make properties satisfied by d_1 more natural than those satisfied by d_2 , so every precisification will make the denotation of 'David Chalmers' fall into the extension of 'conscious'.

We can say the same thing about Alice's original statement: *That is where Harold and I first danced the rumba*. Since one can't *first* dance the rumba with Harold in two different places, it seems Alice's statement can't be true relative to more than one precisification of 'That'. But really the phrase after 'is' is also vague, and there is a penumbral connection (via naturalness) between it and the demonstrative. Hence we can say Alice's statement is supertrue without appealing to any *mysterious* penumbral connections.

3. *McGee and McLaughlin's Challenge*

Vann McGee and Brian McLaughlin (2000) raise a challenge for supervenient approaches to the Problem of the Many that uses belief reports in much the way that Schiffer's problem uses speech reports. They fear that without further development, the supervenient theory cannot distinguish between the *de re* and *de dicto* readings of (4).

(4) Ralph believes that there is a snow-capped mountain within sight of the equator.

They claim, correctly, that (4) should have both a *de dicto* reading and a *de re* reading, where in the latter case it is a belief about Kilimanjaro. The problem with the latter case is unclear how Ralph's belief can be about Kilimanjaro itself. To press the point, they consider an atom at or around the base of Kilimanjaro, called Sparky, and define "Kilimanjaro(+)" to be the body of land constituted ... by the atoms that make up Kilimanjaro together with Sparky [and] Kilimanjaro(-) [to] be the body of land constituted ... by the atoms that make up Kilimanjaro other than Sparky." (129) The problem with taking (4) to be true on a *de re* reading is that "there isn't anything, either in his mental state or in his neural state or in his causal relations with his environment that would make one of Kilimanjaro(+) and Kilimanjaro(-), rather than the other, the thing that Ralph's belief is about." (146) So if the truth of (4) on a *de re* reading requires that Ralph believes a singular, or object-dependent, proposition, about one of Kilimanjaro(+) and Kilimanjaro(-), then (4) cannot be true. Even worse, if the truth of (4) requires that Ralph both that Ralph believes a singular proposition about Kilimanjaro(+), that it is a snow-capped mountain within sight of the equator, and the same proposition about Kilimanjaro(-), then given some knowledge about mountains on Ralph's part, (4) cannot be true, because that would require Ralph to mistakenly believe there are two mountains located roughly where Kilimanjaro is located.

We should not be so easily dissuaded. It is hard to identify exactly which features of Ralph's "mental state or neural state or causal relations with his environment" that make it the case that he believes that two plus two equals four, but does not believe that two quus two equals four. (I assume Ralph is no philosopher, so lacks the concept QUUS.) I doubt, for example, that the concept PLUS has some causal influence over Ralph that the concept QUUS lacks. But Ralph does have the belief involving PLUS, and not the belief involving QUUS. He has this belief not merely in virtue of his mental or neural states, or his causal interactions with his environment, but in

virtue of the fact that PLUS is a more natural concept than QUUS, and hence is more eligible to be a constituent of his belief.

So if Kilimanjaro(+) is more natural than Kilimanjaro(-), it will be a constituent of Ralph's belief, despite the fact that there is no other reason to say his belief is about one rather than the other. Now, in reality Kilimanjaro(+) is no more natural than Kilimanjaro(-). But according to any precisification, one of them will be more natural than the other, for precisifications determine content by determining relative naturalness. Hence if Ralph has a belief with the right structure, in particular a belief with a place for an object (roughly, Kilimanjaro) and the property *being within sight of the equator*, then on every precisification he has a singular belief that a Kilimanjaro-like mountain is within sight of the equator. And notice that since naturalness determines both mental content and verbal content, on every precisification the constituent of that belief will be the referent of 'Kilimanjaro'. So even on a *de re* reading, (4) will be true.

Schiffer's problem showed that we should not take precisifications to be defined merely over single sentences. McGee and McLaughlin's problem shows that we should take precisifications to set the content not just of sentences, but of mental states as well. Precisifications do not just assign precise content to every contentful linguistic token, but to every contentful entity in the world, including beliefs. This makes the issue of penumbral connections that we discussed in section two rather pressing. We already noted the need to establish penumbral connections between separate uses of demonstratives. Now we must establish penumbral connections between words and beliefs. The idea that precisifications determine content by determining relative naturalness establishes these connections.

To sum up, McGee and McLaughlin raise three related problems concerning *de re* belief. Two of these concern belief reports. First, how can we distinguish between *de re* and *de dicto* reports? If I am right, we can distinguish between these just the way Russell suggested, by specifying the scope of the quantifiers. McGee and McLaughlin suspect this will not work because in general we cannot argue from (5) to (6), given the vagueness of 'Kilimanjaro'.

- (5) Kilimanjaro is such that Ralph believes it to be within sight of the equator.
- (6) There is a mountain such that Ralph believes it to be within sight of the equator.

Whether or not we want to accept a semantics in which we must restrict existential generalisation in this way as a general rule, we can give an independent argument that (6) is true whenever (4) is true on a *de re* reading (i.e. whenever (5) is true). The argument is just that on every precisification, the subject of Ralph's salient singular belief is a mountain, so (6) is true on every precisification. This argument assumes that there is a penumbral connection between the subject of this belief, as we might say the referent of 'Kilimanjaro' in his language of thought³, and the word 'mountain'. But since we have already established that there is such a connection between 'Kilimanjaro' in his language of thought and 'Kilimanjaro' in public language, and there is obviously a connection between 'Kilimanjaro' in public language and the word 'mountain', as 'Kilimanjaro is a mountain' is supertrue, this assumption is safe. So the second puzzle McGee and McLaughlin raise, how it can be that the relevant *de re* reports can be *true*, has also been addressed.

There is a third puzzle McGee and McLaughlin raise that the reader might think I have not addressed. How can it be that Ralph can actually have a *de re* belief concerning Kilimanjaro? I have so far concentrated on belief reports, not merely on beliefs, and my theory has relied crucially on correlations between the vagueness in these reports and the vagueness in the underlying belief. It might be thought that I have excluded the most interesting case, the one where Ralph has a particular belief with Kilimanjaro itself as a constituent. While I will end up denying Ralph can have such a belief, I doubt this a problematic feature of my view. The theory outlined here denies that Ralph has object-dependent beliefs, but not that he has *de re* beliefs. I deny that Ralph has a belief that has Kilimanjaro(+) as a constituent, but it is hard to see how Ralph *could* have such a belief, since it very hard to see how he could have had a belief that has Kilimanjaro(+) rather than Kilimanjaro(-) as its subject. (This was McGee and McLaughlin's fundamental point.) If we think that having a *de re* belief implies having a belief whose content is an object-dependent proposition, then we must deny that there are *de re* beliefs about Kilimanjaro. Since there is no object that is determinately a constituent of the proposition Ralph believes, it is a little hard to maintain that he believes an object-dependent proposition.⁴ But this is not the only way to make sense of *de re* beliefs.

³ I do not mean here to commit myself to anything like the language of thought hypothesis. This is just being used as a convenient shorthand.

⁴ This is hard, but not perhaps impossible. One might say that on every precisification, Ralph believes a proposition that has a mountain as a constituent, and hence as an essential part.

Robin Jeshion has argued that whether a belief is *de re* depends essentially on its role in cognition. “What distinguishes *de re* thought is its structural or organisational role in thought” (Jeshion 2001: 67).⁵ I won’t rehearse Jeshion’s arguments here, just their more interesting conclusions. We can have *de re* beliefs about an object iff we have a certain kind of mental file folder for the object. This folder need not be generated by acquaintance with the object, so acquaintanceless *de re* belief is possible. Indeed, the folder could have been created defectively, so there is no object that the information in the folder is about.⁶ In this case, the contents of the folder are subjectless *de re* beliefs. Jeshion doesn’t discuss this, but presumably the folder must not have been created purely to be the repository for information about the bearer of a certain property, whoever or whatever that is. We have to rule out this option if we follow Szabó (2000) in thinking the folder metaphor plays a crucial role in explaining our talk and thought involving descriptions. Provided the folder was created with the intent that it record information about some object, rather than merely information about whatever object has a particular property, its contents are *de re* beliefs. (To allow for distinct folders ‘about’ non-existent objects, we must allow that it is possible that such folders *do* have their reference fixed by their contents, but as long as this was not the intent in creation these folders can suffice for *de re* belief. This point lets us distinguish between my folder for *Vulcan* and my folder for *The planet causing the perturbations of Mercury*. Both are individuated by the fact that they contain the proposition *This causes the perturbations of Mercury*. It is this feature of the folder that fixes their reference, or in this case their non-reference. Only in the latter case, however, was this the intent in creating the folder, so its contents are *de dicto* beliefs, while the contents of the former are *de re* beliefs.)

Now we have the resources to show how Ralph can have *de re* beliefs concerning Kilimanjaro. When Ralph hears about it, or sees it, he opens a file folder for Kilimanjaro. This is not intended to merely be a folder for the mountain he just heard about, or saw. It is intended to be a folder for *that*. (Imagine here that I am demonstrating the mountain in question.) The Kripkenstein point about referential indeterminacy applies to folders as much as to words. This point is closely related to Kripke’s insistence that his indeterminacy argument does not rely on behaviourism. So if Ralph’s folder is to have a reference, it must be fixed in part by the

⁵ I don’t know if Jeshion would accept the corollary that if belief is too unstructured to allow for the possibility of such organisational roles, then there is no *de re* belief, but I do.

⁶ Which is not just to say that there is no object that has all the properties in the folder. This is neither necessary nor sufficient for the folder to be about the object, as Kripke’s discussion of ‘famous deeds’ descriptivism should make clear.

naturalness of various putative referents. But that is consistent with Ralph's folder containing *de re* beliefs, since unless Ralph is a certain odd kind of philosopher, he will not have in his folder that Kilimanjaro is peculiarly eligible to be a referent. So the referent of the folder is not fixed by its contents (as the referent for a folder about *The mountain over there, whatever it is*, would be, or how the referent for a folder about *The natural object over there, whatever it is*, would be), and the contents of this folder are still *de re* beliefs Ralph has about Kilimanjaro. This was a bit roundabout, but we have seen that the Problem of the Many threatens neither the possibility that Ralph is the subject of true *de re* belief ascriptions, nor that he actually has *de re* beliefs.

4. *Vague Objects*

"I think the principle that to be is to be determinate is *a priori*, and hence that it is *a priori* that there is no *de re* vagueness". (Jackson 2001: 657-8)

So do I. I also think there are a few arguments for this claim, though some of them may seem question-begging to the determined defender of indeterminate objects. Most of these arguments I will just mention, since I assume the reader has little desire to see them detailed *again*. One argument is just that it is obvious that there is no *de re* vagueness. Such 'arguments' are not worthless. The best argument that there are no true contradictions is of just this form, as Priest (1998) shows. And it's a good argument! Secondly, Russell's point that *most* arguments for *de re* vagueness involve confusing what is represented with its representation still seems fair (Russell 1923). Thirdly, even though the literature on this is a rather large, it still looks like the Evans-Salmon argument against vague identities works, at least under the interpretation David Lewis gives it, and this makes it hard to see how there could be vague objects (Evans 1978, Salmon 1981, Lewis 1988). Fourthly, Mark Heller (1996) argues that we have to allow that referential terms are semantically vague. He says we have to do so to explain context dependence but there are a few other explanatory projects that would do just as well. Since semantic conceptions of vagueness can explain all the data that are commonly taken to support ontological vagueness, it seems theoretically unparsimonious to postulate ontological vagueness too. That's probably enough, but let me add one more argument to the mix. Accepting that Kilimanjaro is be a vague material object distinct from both Kilimanjaro(+) and Kilimanjaro(-) has either metaphysical or logical costs. To prove this, I derive some rather unpleasant metaphysical conclusions from the assumption that Kilimanjaro is vague. The proofs will use some contentious principles of classical logic, but rejecting those, and hence rejecting classical logic, would be a substantial logical

cost. The most contentious such principle used will be an instance of excluded middle: Sparky is or is not a part of Kilimanjaro. I also assume that if for all x other than Sparky that x is a part of y iff it is a part of z , then if Sparky is part of both y and z , or part of neither y nor z , then y and z coincide. If someone can contrive a mereological theory that rejects this principle, it will be immune to these arguments.

It is very plausible that material objects are individuated by the materials from which they are composed, so any coincident material objects are identical. Properly understood, that is a good account of what it is to be material. The problem is getting a proper understanding. Sider (1997) interprets it as saying that no two non-identical material objects coincide *right now*. His project ends up running aground over concerns about sentences involving counting, but his project, of finding a strong interpretation of the principle is intuitively compelling. David Lewis (1986: Ch. 4) defends a slight weaker version: no two non-identical material objects coincide *at all times*. Call this the *strong composition principle* (scp). The scp is (classically) inconsistent with the hypothesis that Kilimanjaro is vague. If Sparky is part of Kilimanjaro, then Kilimanjaro and Kilimanjaro(+) always coincide. If Sparky is not part of Kilimanjaro then Kilimanjaro and Kilimanjaro(-) always coincide. Either way, two non-identical objects always coincide, which the scp does not allow.

Some think the scp is refuted by Gibbard's example of Lump and Goliath (Gibbard 1975). The most natural response to Gibbard's example is to weaken our individuation principle again, this time to: no two non-identical material objects coincide *in all worlds at all times*. Call this the *weak compositional principle* (wcp). Since there are worlds in which Goliath is composed of bronze, but Lump is still a lump of clay in those worlds, Lump and Goliath do not refute the wcp. Some may think that even the wcp is too strong⁷, but most would agree that if vague objects violated the wcp, that would be a reason to believe they don't exist.

Given a plausible metaphysical principle, which I call Crossover, vague objects will violate the wcp. As shown above, Kilimanjaro actually (always) coincides with Kilimanjaro(+) or Kilimanjaro(-), but is not identical with either. Crossover is the following principle:

Crossover: For any actual material objects x and y there is an object z that coincides with x in the actual world and y in all other worlds.

⁷ Kit Fine (1994) does exactly this.

Given that arbitrary fusions exist, Crossover is entailed by, but does not entail, the doctrine of arbitrary modal parts: that for any object o and world w , if o exists in w then o has a part that only exists in w . But Crossover does not have the most surprising consequence of the doctrine of arbitrary modal parts: that for any object o there is an object that has essentially all the properties o actually has.

Let K1 be the object that coincides with Kilimanjaro in this world and Kilimanjaro(+) in all other worlds. Let K2 be the object that coincides with Kilimanjaro in this world and Kilimanjaro(-) in all other worlds. If Sparky is part of Kilimanjaro then K1 and Kilimanjaro(+) coincide in all worlds, but they are not identical, since it is determinate that Sparky is actually part of Kilimanjaro(+) and not determinate that it is part of K1. If Sparky is not part of Kilimanjaro then K2 and Kilimanjaro(-) coincide in all worlds, but they are not identical, since it is determinate that Sparky is not actually part of Kilimanjaro(-) and not determinate that it is not part of K2. Either way, we have a violation of the wcp. So the following three claims are (classically) inconsistent.

- (a) Crossover.
- (b) The wcp.
- (c) Kilimanjaro is a vague object that indeterminately has Sparky as a part.

I think the first two are highly plausible, so accepting (c) is costly. I already noted the plausibility of the wcp, so the focus should be on Crossover. On Lewis's account of modality, it is clearly true, as is the stronger doctrine of arbitrary modal parts. On a fictionalist theory of modality based on Lewis's account, it is still true, or at least true in the fiction that we must adopt to make sense of modal talk. So the principle is not without merits. And dialectically, opposing Crossover will be problematic for the believer in vague objects. Either an object's modal profile is determined by its categorical properties or it isn't. If it is, then the wcp will entail the scp, so by the above reasoning vague objects will be inconsistent with the wcp. If it is not, then it is hard to see why an object could not have a completely arbitrary modal profile, say the profile of some other ordinary material object. But that means Crossover is true, and again we cannot have both the wcp and vague objects. Probably the best way out for the believer in vague objects will be to short-circuit this reasoning by abandoning classical logic, presumably by declining to endorse the version of excluded middle with which I started. But that is undoubtedly a costly move, particularly for a supervaluationist.

5. *McKinnon on Coins and Precisifications*

Most of our discussions of the Problem of the Many relate to the vagueness in a single singular term, and a single ordinary object. As McKinnon reminds us, however, there is not just one mountain in the world, there are many of them, and supervaluationists are obliged to say plausible things about statements that are about many mountains. Or, to focus on McKinnon's example, we must not only have a plausible theory of coins, but of coin exhibitions. These do raise distinctive problems. Imagine we have an exhibition with, as we would ordinarily say, 2547 coins, each numbered in the catalogue. So to each number n there correspond millions of coin-like entities, coin*s in Sider's helpful phrase (Sider 2001), and each precisification assigns a coin* to a number. In general, Sider holds that something is an F^* iff it has all the properties necessary and sufficient for being an F except the property of not massively overlapping another F . There are some interesting questions about how independent these assignments can be. If one precisification assigns coin* c_1 to n_1 , and another assigns coin* c_2 to n_2 (distinct from n_1) then is there a guaranteed to be a precisification that assigns both c_1 to n_1 and c_2 to n_2 ? In other words, may the precisifications of each numeral (construed as a coin denotation) be *independent* of each other? The following example suggests not. Say C_j is the set of coin*s that are possible precisifications of j . This set may be vague because of higher-order vagueness, but set those difficulties aside. If every member of C_{1728} has a duplicate in C_{1729} , then presumably only precisifications that assigned duplicates to '1728' and '1729' would be admissible. If the exhibition has two Edward I pennies on display to show the obverse and reverse, and miraculously these coins are duplicates, such a situation will arise.

This case is fanciful, so we don't know whether in reality the precisifications of the numerals are independent. We probably can't answer this question, but this is no major concern. McKinnon has found a question which the supervaluationist should feel a need to answer, but to which neither answer seems appropriate. Say that a precisification is *principled* iff there is some not-too-disjunctive property F such that for each numeral n , the precisification assigns to n the F -est coin* in C_n . If F does not come in degrees, then the precisification assigns to n the F in C_n . McKinnon's question to the supervaluationist is: Are all precisifications principled? He aims to show either answering 'yes' or 'no' gets the supervaluationist in trouble. 'Yes' leads to there being too few precisifications; 'No' leads to there being too many. Let us look at these in order.

I have little to say for now on the first horn of this dilemma. McKinnon's survey of principled precisifications only considers cases where F is intrinsic, and I postpone for now investigation of extrinsic

principles. Nevertheless, he does show that if F must be intrinsic, then there are not enough principled precisifications to generate all the indeterminacy our coin exhibit intuitively displays. The other horn is trickier.

A precisification must not only assign a plausible coin* to each numeral, it must do so in such a way that respects penumbral connections. McKinnon thinks that unprincipled, or *arbitrary* precisifications, will violate (NAD) and (NAS).

Non-Arbitrary Differences (NAD): For any coin and non-coin, there is a principled difference between them which forms the basis for one being a coin and the other being a non-coin.

Non-Arbitrary Similarities (NAS): For any pair of coins, there is a principled similarity between them which forms the basis for their both being coins.

McKinnon holds these are true, so they should be true on all precisifications, but they are not true on unprincipled precisifications, so unprincipled precisifications are unacceptable. The motivation for (NAD) and (NAS) is clear. When we list the fundamental properties of the universe, we will not include *being a coin*. Coinness doesn't go that deep. So if some things are coins, they must be so in virtue of their other properties. From this (NAD) and (NAS) follow.

The last step looks dubious. Consider any coin, for definiteness say the referent of '1728', and a coin* that massively overlaps it. The coin* is not a coin, so (a) one of these is a coin and the other is not, and (b) the minute differences between them cannot form the basis for a distinction between coins and non-coins. Hence (NAD) and (NAS) fail. At best, it seems, we can justify the following claims. If something is a coin* and something else is not, *then* there is a principled difference between them that makes one of them a coin* and the other not. Something is a coin iff it is a coin* that does not excessively overlap a coin. If this is the best we can do at defining 'coin', then the prospects for a *reductive* physicalism about coins might look a little dim, though this is no threat to a physicalism about coins that stays neutral on the question of reduction. (I trust no reader is an anti-physicalist about coins, but it is worth noting how vexing questions of reduction can get even when questions of physicalism are settled.)

So I think this example refutes (NAD) and (NAS). Do I beg some questions here? Well, my counterexample turns crucially on the existence of kinds of objects, massively overlapping coin*s, that some

people reject, and indeed that some find the most objectionable aspect of the supervaluationist solution. But this gets the burden of proof the wrong way around. I was not trying to refute (NAD) and (NAS). I just aimed to parry an argument based on those principles. I am allowed to appeal to aspects of my theory in doing so without begging questions. I do not want to rest too much weight on this point, however, for issues to do with who bears the burden of proof are rarely easily resolved, so let us move on.

My main response to McKinnon's dilemma is another dilemma. If the principled similarities and differences in (NAD) and (NAS) must be intrinsic properties, then those principles are false, because there is no principled intrinsic difference between a coin and a token, or a coin and a medal. If the principled similarities and differences in (NAD) and (NAS) may be extrinsic properties, then those principles may be true, but then the argument that there are not enough principled precisifications fail, since now we must consider precisifications based on extrinsic principles. Let's look at the two halves of that dilemma in detail, in order.

A subway token is not a coin. Nor is a medal.⁸ But in their intrinsic respects, subway tokens often resemble certain coins more than some coins resemble other coins. Imagine we had a Boston subway token (which looks a bit like an American penny, but larger), an American penny, a British 20p piece (which is roughly heptagonal) and an early Australian holey dollar (which has a hole in it). There is no non-disjunctive classification of these by intrinsic properties that includes the penny, the 20p piece and the holey dollar in one group, and the subway token in the other. Any group that includes the penny and the other coins will include the token as well. So if we restrict attention to intrinsic similarities and differences, (NAD) and (NAS) are false.

There is a difference between these coins and the subway token. The coins were produced with the intent of being legal tender, the token was not. Perhaps we can find a difference between coins and non-coins based on the intent of their creator.⁹ This might make (NAD) and (NAS) true. But note that given the theory of precisifications developed in section 3, on every precisification, one and only one of the precisifications of '1728'

⁸ Some people I have asked think tokens are coins, but no one thinks medals are coins, so if you (mistakenly) think tokens are coins, imagine all my subsequent arguments are phrased using medals rather than tokens.

⁹ Note that I say little here about what the intent of the creator must be. I don't think that the intent must always be to create legal tender. A ceremonial coin that is created, for example, to be tossed before the start of a sporting match is still a coin, although it is not intended to be tender. But intent still matters. If someone had made a duplicate of that ceremonial coin with the intent of awarding it as a medal to the victorious captain, it would be a medal and not a coin.

will be the subject of an intention on the part of its manufacturer. Just which of the objects is the subject of this intent will vary from precisification to precisification, but there is only one on every precisification. So we can say that on every precisification, *the* coin is the one where the intent of its creator was that it be used in a certain way. Indeed, on any precisification we may have antecedently thought to have existed, we can show that precisification to be principled by taking *F* to be the property *being created with intent of being used in a coin-like way*.¹⁰ So now we can say that restricting attention to the principled precisifications does not unduly delimit the class of precisifications.

Let's sum up. To argue against the possibility of unprincipled precisifications, McKinnon needed to justify (NAD) and (NAS). But these are only true when we allow 'principled differences' to include differences in creatorial intent. And if we do that we can see that every *prima facie* admissible precisification is principled, so we can give an affirmative answer to McKinnon's question.

It might be objected that this move relies heavily on the fact that for many artefacts creative intent is constitutive of being the kind of thing that it is. But a Problem of the Many does not arise only for artefacts, so my solution does not generalise. This is little reason for concern since McKinnon's problem does not generalise either. (NAD) and (NAS) are clearly false when we substitute 'mountain' for 'coin'. Consider a fairly typical case where it is indeterminate whether we have one mountain or two.¹¹ In this case it might be not clear whether, for example, we have one mountain with a southern and a northern peak, or two mountains, one of them a little north of the other. Whether there is one mountain here or two, clearly the two peaks exist, and their fusion exists too. The real question is which of these three things is a mountain. However this question is resolved, a substitution instance of (NAD) with the two objects being the southern peak and the fusion of the two peaks will be false. So in this case a relatively unprincipled precisification will be acceptable. The point here is that mountain*s that are not mountains exist (either the peaks or their fusion will do as examples), and that suffices to refute McKinnon's alleged penumbral connections and allow, in this case, a negative answer to his question.

¹⁰ Because of the problems raised in the previous footnote, I will not try and say just what this intention amounts to. There are complications when (a) the creator is a corporate entity rather than an individual and (b) the coins are mass-produced rather than produced individually. But since the story is essentially the same, I leave the gruesome details out here.

¹¹ This case is rather important in the history of the problem, because its discussion in Quine (1960) is one of the earliest presentations in print of anything like the problem of the many.

6. *Sorensen on Direct Reference*

According to orthodoxy, we can use descriptions to determine the reference of names without those descriptions becoming part of the meaning of the name. This, apparently, is what happened when Leverrier introduced 'Neptune' to *name*, not merely describe, the planet causing certain perturbations, and when someone introduced 'Jack the Ripper' to *name*, not merely describe, the person performing certain murders. So let us introduce 'Acme' as the name for the first tributary of the river Enigma. As Sorensen suggests, this can create certain problems.

When [explorers] first travel up the river Enigma they finally reach the first pair of river branches. They name one branch 'Sumo' and the other 'Wilt'. Sumo is shorter but more voluminous than Wilt. This makes Sumo and Wilt borderline cases of 'tributary' ... 'Acme' definitely refers to something, even though it is vague whether it refers to Sumo and vague whether it refers to Wilt. (Sorensen 2000: 180).

If 'Acme', 'Sumo' and 'Wilt' are all vague names related in this way, Sorensen thinks the supervenient has a problem. The sentences 'Acme is Sumo' and 'Acme is Wilt' both express propositions of the form $\langle x = y \rangle$. For exactly one of them, x is y . Since the proposition contains just the objects x and y (and the identity relation) but not their route into the proposition, there is no vagueness in the proposition. Hence there is no way to precisify either proposition. So a supervenient cannot explain how these propositions are vague.

This is no problem for supervenientism, since supervenientism says that *sentences*, not *propositions*, are vague. Indeed, most supervenientists would say that no proposition is ever vague. Thinking they are vague is just another instance of the fallacy Russell identified: attributing properties of the representation to the entity, in this case a proposition, represented.

But maybe there is a problem in the area. One natural way of spelling out the idea that names directly refer to objects is to say that the *meaning* of a name is its referent. And one quite plausible principle about precisifications is that precisifications must not *change* the meaning of a term, they may merely provide a meaning where none exists. Now the supervenient has a problem. For it is true that one of Sorensen's identity sentences is true in virtue of its meaning, since its meaning determines that it expresses a proposition of the form $\langle x = x \rangle$. But each sentence is false on some precisifications, so some precisifications change the meaning of the terms involved.

The best way to respond to this objection is simply to bite the bullet. We can accept that some precisifications alter meanings provided we can provide some other criteria for acceptability of precisifications. I offered one such proposal in section 2. An acceptable precisification takes the partial order *more natural than*, turns it into a complete order without changing any of the relations that already exist, and uses this new relation to generate meanings. If we proceed in this way it is possible, for all we have hitherto said, that on every precisification the proposition expressed by 'Acme is Sumo' will be of the form $\langle x = y \rangle$, so just the named object, rather than the method of naming, gets into the proposition. The central point is that since precisifications apply to the processes that turn semantic intentions into meanings, rather than to sentences with meanings, there is no guarantee they will preserve meanings. But if we like directly referential theories of names we should think this perfectly natural. If names are directly referential then Sorensen's argument that there are vague sentences that are true in virtue of their meaning works. But this is consistent with supervenience.

One challenge remains. If precisifications change meanings, why should we care about them, or about what is true on all of them? This is not a new challenge; it is a central plank in Jerry Fodor and Ernest Lepore's (1996) attack on supervenience. A simple response is just to say that we should care about precisifications because this method delivers the right results in all core cases, and an intuitively plausible set of results in contentious cases. This kind of instrumentalism about the foundations of a theory is not always satisfying.¹² But if that's the biggest problem supervenience has, they should be able to sleep a lot easier than the rest of us.

7. *Conclusions and Confessions*

I have spent a fair bit of time arguing that supervenience is not vulnerable to a few challenges based on the Problem of the Many. Despite doing all this, I don't believe supervenience to be quite true. So why spend this time? Because the true theory of vagueness will be a *classical semantic* theory, and everything I say about supervenience above applies *mutatis mutandis* to all classical semantic theories. I focussed on supervenience because it is more familiar and more popular, but I need not have.

¹² The largest debate in the history of philosophy of economics concerned whether we could, or should, be instrumentalists about the ideally rational agents at the core of mainstream microeconomic theory. See Friedman (1953) for the classic statement of the instrumentalist position, and Hausman (1992) for the most amusing and enlightening of the countably many responses.

What is a classical semantic theory? That's easy - it's a theory that is both classical and semantic. What is a classical theory? It is one that incorporates vagueness while preserving classical logic. How much of classical logic must we preserve? That's a hard question, though it is relevant to determining whether supervaluationism is (as it is often advertised) a classical theory. Williamson (1994) notes that supervaluationism does not preserve classical inference rules, and Hyde (1997) notes that it does not preserve some classically valid multiple-conclusion sequents. Keefe (2000) argues that neither of these constitutes an important deviation from classical logic. I'm inclined to disagree with Keefe on both points. Following Read (2000), I take it that the best response to the anti-classical arguments in Dummett (1991) takes the essential features of classical logic to be its inferential rules as formulated in a multiple-conclusion logic. But we need not adjudicate this dispute here. Why should we want a classical theory? The usual arguments for it are based on epistemic conservatism, and I think these arguments are fairly compelling. I also think that no non-classical theory will be able to provide a plausible account of quantification.¹³

What is a semantic theory? It is one that makes vagueness a semantic phenomenon. It is not necessarily one that makes vagueness a *linguistic* phenomenon. That would be absurd in any case, since clearly some non-linguistic entities, maps, beliefs and pictures for example, are vague. But the more general idea that vagueness is a property only of representations is quite attractive. It links up well with the theory of content Lewis outlines in "Languages and Language" - all Languages (in his technical sense) are precise, vagueness in natural language is a result of indecision about which Language we are speaking.

Trenton Merricks (2001) argues against this picture, claiming that all semantic vagueness (he says 'linguistic', but ignore that) must arise because of metaphysical or epistemic vagueness. He claims that if (17) is vague, then so is (18), and (18)'s vagueness must be either metaphysical or semantic.

(17) Harry is bald.

(18) 'Bald' describes Harry.

One might question the inference from (17)'s vagueness to (18) - on some supervaluational theories if (17) is vague then (18) is false. But I will let that pass, for there is a simpler problem in the argument. Merricks claims that if (18)

¹³ See the last section of my "True, Truer, Truest" (manuscript) for a detailed defence of this claim.

is vague, then it is vague whether 'Bald' has the property *describing Harry*, and this is a kind of metaphysical vagueness. It is hard to see how this follows. If there is metaphysical vagueness, there is presumably some object *o* and some property *F* such that it is vague whether the object has the property. Presumably the object here is the word 'bald' and the property is *describing Harry*. But words alone do not have properties like *describing Harry*. At best, words in languages do so. So maybe the object can be the ordered pair <'Bald', *l*>, where *l* is a language. But which one? Not one of Lewis's Languages, for then it is determinate whether <'Bald', *l*> has the property *describing Harry*. So maybe a natural language, perhaps English! But it is doubly unclear that English is an object. First, it is unclear whether we should reify natural languages to such a degree that we accept that 'English' refers to anything at all. Secondly, if we say 'English' does refer, why not say that it refers to one of Lewis's Languages, though it is vague which one? That way we can say that the sentence '*Bald* in English describes Harry' is vague without there being any object that vaguely instantiates a property. Now on a supervenient theory this approach may have the unwanted consequence that "English is a precise language" is true, since it is true on all precisifications. It does not seem that *this* problem for the supervenient generalises to be a problem for all semantic theories of vagueness, so Merricks has raised no general problem for semantic theories of vagueness. (The problem for the supervenient here is not new. For some discussion see Lewis's response, in "Many, but Almost One" to the objection, there attributed to Kripke, that the supervenient account makes it true that all words are precise.)

If we have a classical semantic theory that provides a concept of determinateness, then we can define acceptable precisifications as maximal consistent extensions of the set of determinate truths. Given that, it follows pretty quickly that determinate truth implies truth on all precisifications. And this is sufficient for the major objections canvassed above to get a foothold, and hence be worthy of response, though as we have seen none of them will ultimately succeed. Still, our theory may differ from supervenientism in many ways. For one thing, it might *explain* determinateness in ways quite different from those in supervenientism. For example, the theory in Field (2000) is a classical semantic theory¹⁴, but it clearly goes beyond supervenient theory because it has an

¹⁴ At least, it strikes me as a classical semantic theory. Ryan Wasserman has tried to convince me that properly understood, it is really an epistemic theory. Space prevents a thorough account of why I think Field's theory is flawed. Briefly, I think the point in Leeds (2000) that Field's concept of a numerical degree of belief needs *substantially* more explanation than Field gives it can be developed into a conclusive refutation.

interesting, if ultimately flawed, explanation of determinateness in terms of Shafer functions. Other classical semantic theories may differ from supervaluationism by providing distinctive theories of higher order vagueness.

The most promising research programs in vagueness are within the classical semantic framework. Like all research programs, these programs need a defensive component, to fend off potential refutations and crisis. This avoids unwanted crises in the program, and as we have seen here we can learn a bit from seeing how to defend against certain attacks. There will undoubtedly be more challenges in the time ahead, but for now the moves in this paper brings the defensive side of the program up to date.

References

- Dummett, Michael (1991) *The Logical Basis of Metaphysics*. Cambridge, MA: Harvard.
- Evans, Gareth (1978) "Can There be Vague Objects?" *Analysis* 38: 208.
- Field, Hartry (2000) "Indeterminacy, Degree of Belief, and Excluded Middle" *Noûs* 34: 1-30.
- Fine, Kit (1994) "Compounds and Aggregates" *Noûs* 28: 137-58.
- Fodor, Jerry and Ernest Lepore (1996) "What Cannot be Valuated Cannot be Valuated, and it Cannot be Supervaluated Either" *Journal of Philosophy* 93: 516-535.
- Friedman, Milton (1953) "The Methodology of Positive Economics." in his *Essays in Positive Economics*, Chicago: University of Chicago Press. pp. 3-43.
- Gibbard, Allan (1975) "Contingent Identity" *Journal of Philosophical Logic* 4: 187-222.
- Hausman, Daniel (1992) "Why Look Under the Hood?" in his *Essays in Philosophy and Economic Methodology*, Cambridge: Cambridge University Press. pp. 70-3.
- Heller, Mark (1996) "Against Metaphysical Vagueness" *Philosophical Perspectives* 10: 177-85.
- Hyde, Dominic (1997) "From Heaps and Gaps to Heaps of Gluts" *Mind* 106: 641-660
- Jackson, Frank (2001) "Responses" *Philosophy and Phenomenological Research* 62: 653-64.
- Jeshion, Robin (2001) "Acquaintanceless *De Re* Belief" in M. O'Rourke *et al* (eds.) *Truth and Meaning: Investigations in Philosophical Semantics*, New York: Seven Bridges Press, pp. 53-74.
- Keefe, Rosanna (2000) *Theories of Vagueness*. Cambridge: Cambridge University Press.
- Kripke, Saul (1979) "A Puzzle About Belief" in A. Margalit (ed.) *Meaning and Use*, Dordrecht: Reidel. pp. 239-83.
- Kripke, Saul (1980) *Naming and Necessity*. Cambridge: Harvard University Press.
- Kripke, Saul (1982) *Wittgenstein on Rules and Private Language*. Oxford: Blackwell.
- Leeds, Stephen (2000) "A Disquotationalist Looks at Vagueness" *Philosophical Topics* 28: 107-28.
- Lewis, David (1975) "Languages and Language" in Keith Gunderson (ed.) *Minnesota Studies in the Philosophy of Science*, vol. 7, Minneapolis: University of Minnesota Press, pp 3-35.
- Lewis, David (1983) "New Work for a Theory of Universals" *Australasian Journal of Philosophy* 61: 343-77.
- Lewis, David (1984) "Putnam's Paradox" *Australasian Journal of Philosophy* 62: 221-36.
- Lewis, David (1986) *On the Plurality of Worlds*. Oxford: Blackwell.
- Lewis, David (1988) "Vague Identity: Evans Misunderstood" *Analysis* 48: 128-30.

- Lewis, David (1993) "Many, but Almost One" in K. Campbell *et al* (eds.) *Ontology, Causality, and Mind: Essays on the Philosophy of D. M. Armstrong* Cambridge: Cambridge University Press, pp. 23-38.
- McGee, Vann and Brian McLaughlin (2000) "The Lessons of the Many" *Philosophical Topics* 28: 129-51.
- McKinnon, Neil (2002) "Supervaluations and the Problem of the Many" *Philosophical Quarterly* xx: xx-xx
- Merricks, Trenton (2001) "Varieties of Vagueness" *Philosophy and Phenomenological Research* 62: 145-57.
- Priest, Graham (1998) "What Is So Bad About Contradictions?" *Journal of Philosophy* 95: 410-426.
- Putnam, Hilary (1973) "Meaning and Reference" *Journal of Philosophy* 70: 699-711.
- Putnam, Hilary (1980) "Models and Reality" *Journal of Symbolic Logic* 45: 464-82.
- Quine, W. V. O. (1960) *Word and Object*. Cambridge, MA: MIT Press.
- Read, Stephen (2000) "Harmony and Autonomy in Classical Logic" *Journal of Philosophical Logic* 29: 123-54.
- Russell, Bertrand (1923) "Vagueness" *Australasian Journal of Philosophy and Psychology* 1: 84-92.
- Salmon, Nathan (1981) *Reference and Essence*. Princeton: Princeton University Press.
- Schiffer, Stephen (1998) "Two Issues of Vagueness" *Monist* 81: 193-214.
- Schiffer, Stephen (2000a) "Vagueness and Partial Belief" *Philosophical Issues* 10: 220-57.
- Schiffer, Stephen (2000b) "Replies" *Philosophical Issues* 10: 320-43.
- Sider, Theodore (1997) "All the World's a Stage" *Australasian Journal of Philosophy*
- Sider, Theodore (2001) "Maximality and Intrinsic Properties" *Philosophy and Phenomenological Research* 63: 357-64.
- Sorensen, Roy (2000) "Direct Reference and Vague Identity" *Philosophical Topics* 28: 175-94.
- Szabó, Zoltan Gendler (2000) "Descriptions and Uniqueness" *Philosophical Studies* 101: 29-57.
- Williamson, Timothy (1994) *Vagueness*. London: Routledge.