

## Craziness and Metasemantics<sup>1</sup>

Consider a crazy interpretation of our utterances that has the virtue of being charitable – most of our utterances come out true<sup>2</sup> – but the vice of being crazy. What makes such an interpretation incorrect? David Lewis (and various philosophers since) have pinned their hopes on an eligibility constraint: interpretations that assign more natural properties to predicates are, other things being equal, better. On this picture, our words have fairly determinate meanings (contra Quine [1960] and Kripkenstein [1982]) and, at a pretty good first pass, it is the twin constraints of charity and eligibility that explain why this is so.<sup>3</sup> In his admirable ‘Eligibility and Inscrutibility’, Robert Williams makes trouble for this package. In §1 below, I shall describe some cases that reinforce Williams’ misgivings. In §2 I note a general problem that afflicts the Lewisian vision. In §3 I offer a few constructive suggestions concerning how the ‘crazy interpretation’ problem should

---

<sup>1</sup> I am grateful to Brian Weatherson, an audience at Cornell and especially to Tamar Gendler and Robert Williams for comments and conversation.

<sup>2</sup> We might imagine charity formulated to give a special weighting to the sentences the speaker is most confident of: I am not going to fuss about this for now. I am also simplifying Lewis’ own views on charity to some extent. See fn. 10 below.

<sup>3</sup> Lewis is not aiming to secure a unique interpretation, only to make the level of indeterminacy tolerable. As a supervaluationist, he is happy to allow for indeterminacy owing to vagueness.

be approached. Finally, in §4, I try to shed some light on the role of the interpreter in metasemantics

§1

Consider the following set of crazy interpretation puzzles, cases where eligibility plus charity, as normally formulated, have no hope of saving the day.

*A Belief worlds.*

Suppose there is a planet where someone a lot like me is in an environment that is, qualitatively, exactly like I believe my environment to be: If I have a friend Frank whom I (rightly or wrongly) believe has just gone on a picnic in Kent, then my Twin has a friend that he calls 'Frank' who has just gone on a picnic in a place he calls 'Kent'.

Consider the following 'crazy' interpretation for my utterances: when I say 'Frank,' I am referring to the person that my Twin calls 'Frank', when I say 'Kent' I am referring to the place my Twin calls 'Kent,' when I say 'I', I am referring to my Twin, and so on.

Predicates, by contrast, are generally interpreted in a non-deviant way: 'red' means red, 'negative charge' means negative charge, and so on.<sup>4</sup> Clearly, the crazy interpretation is even more charitable than the standard one. And the crazy interpretation comes out no

---

<sup>4</sup> Of course there will be exceptions for predicates that are constitutively connected to singular terms: 'Arabian' will be reinterpreted so as to coordinate with the crazy interpretation of 'Arabia'.

worse on the score of eligibility: on either interpretation, my predicates score the same. So neither charity nor eligibility can explain the craziness of the interpretation.

If modal realism is true, the problem gets worse. For any subject in a given world, there are crazy interpretations according to which it is the entities in some other possible world that serve as the referents for her singular terms, the domain for her quantifiers, and the anchor for her use of “actually,” etc. As above, it is easy to see that if the world is well chosen, the crazy interpretation will do better on the score of charity than the intended one and no worse on the score of eligibility.

Consider, similarly, the many worlds interpretation of quantum mechanics according to which the possible futures for a system described by its wave function correspond to branching futures that are created by fission. Supposing that vision of reality is correct, there will always be charitable, eligible but crazy interpretations that reckon an individual on one branch to be talking about goings on in another branch.

### *B Mirror Worlds*

Suppose I happen to live in a mirror universe, where there is symmetry in the mosaic of events with respect to some spatial axis. A crazy interpretation will assign mirror Kent as the referent of ‘Kent’ in my mouth, mirror Tom Jones as the referent of ‘Tom Jones’ and so on. Such an interpretation does no worse on considerations of charity and eligibility than its converse.

### *C Micro worlds.*

Suppose it so happens that some micro corner of the actual world is such that there is a mapping from the singular terms I use to microscopic objects in that corner of the world and from predicates that I use to rather natural properties of those microscopic objects such that relative to that interpretation, much of what I say comes out true. Here again, interpreting my utterances as referring to these microentities does no worse on considerations of charity and eligibility than its saner cousin.

### *D. Implications*

One possible response to the worries raised by these cases would be to concede that there is no fact of the matter as to what the individuals in the problematic scenarios are referring to. That concession is compatible with the claim that *our* words have fairly determinate meanings. After all, our world might not fit any of the problematic descriptions. Obviously, this response is not available to the modal realist or the quantum fission lover. But in any case it is deeply unsatisfying: it seems clear that the presence of a mirror universe, belief worlds, and so on, would constitute *no serious obstacle at all* to an individual's capacity to refer to himself and those around him (modulo mild and familiar levels of vagueness)<sup>5</sup>. The crazy interpretations are nowhere near to being acceptable at those worlds. Hence Lewis' constraints of charity plus eligibility need to be supplemented.

---

<sup>5</sup> I take no stand for or against epistemicist accounts of vagueness here.

\$2 There is a further general problem for the charity plus eligibility picture (one that is easily overlooked when one's guiding picture is standard model theory for artificial languages). In formulating Lewis' account, its proponents typically ignore the widespread context dependence that afflicts nearly all predicates of natural language. Consider 'global folk theory' -- 'the sum of platitudes gathered from every walk of life'.<sup>6</sup> On the standard Lewisian picture, charity and eligibility combine to fix the meaning of that theory, though it operates on the naïve assumption that a predicate, as it occurs in one platitude in one context, expresses the same property as when it occurs in another platitude at another context. But this assumption is untenable: many predicates are ambiguous and many others are context-dependent.

Once the assumption of semantic stability is dropped, it is transparently hopeless to assume that charity and eligibility will fix meaning within reasonably determinate limits. For if there is no requirement to interpret tokens of a given predicate in the same way, a plenitude of crazy but acceptable interpretations can easily be concocted. Take a simple example. Suppose someone says ' $3 + 2 = 6$ ' on one occasion. One candidate interpretation is this: on that occasion of use '+' meant multiplication, even though on other occasions of use that symbol expresses addition. (That the speaker will disavow ' $3 + 2 = 6$ ' a little later is perfectly compatible with that hypothesis, since the later occasion may be one where '+' does mean addition.) Charity plus eligibility cannot explain why such a proposal is crazy. More subtle versions of the problem arise when one asks what

---

<sup>6</sup> See Williams, this journal issue, p.

constitutes the content of a shifty predicate – say ‘flat’ or ‘a lot’ -- on a particular occasion of use, given that there is no semantic uniformity across contexts. Note that in the context of the current inquiry, it is not enough to say that the intentions of the speaker fix the meaning, since what we are after is a foundational account of intentionality that offers a sketch of what constitutes the content of intentions.<sup>7</sup>

It might be hoped that this problem arises only when we focus our attention on public language. Perhaps that there is a language of thought – mentalese – in which semantic context dependence is restricted to a few simple indexicals and demonstratives. I am dubious about basing the foundations of semantics on so optimistic an empirical hypothesis. Moreover, if mentalese exists and its predicates are not context-dependent, one would like to know what constitutes this fact. The eligibility plus charity approach, as it is standardly articulated, cannot explain this since it simply takes semantic stability for granted.

§3 If the eligibility plus charity picture fails to give us even the rough outlines of an adequate solution to the crazy interpretation problem, then how is the situation to be rectified? Here, briefly, are six ideas.

---

<sup>7</sup> A more modest version of the eligibility approach assumes a range of facts concerning coarse-grained intentions – with contents individuated as sets of worlds -- and asks after the content of language in that setting. While that is the intent of the relevant segments of Lewis [1969], it is not the spirit in which the eligibility proposal has been discussed and with good reason: eligibility is supposed to explain how our thoughts get magnetized to certain sets of worlds/intensions over others.

(1) Lewis allowed that causal predicates may figure in the reference fixing descriptions for certain singular terms. But he did not expect causal constraints to play a constitutive role in solving the crazy interpretation problem.<sup>8</sup> In so doing, he overlooked a promising resource.

While functionalist dreams of *defining* content in terms of some profile of causal input-output relations were overambitious, causal relations must nonetheless play a central role in foundational accounts of the mind.

Consider *perceptual attention*, the anchor for the most basic deictic vehicles of thought and talk. It is clear that causal constraints can help here in filtering out certain crazy hypotheses. While it would be tendentious to suppose that there is a general causal constraint on semantic reference, it is far more plausible to think that there is a constitutive causal constraint on perception.<sup>9</sup> The analogue of certain crazy interpretation problems for perceptual attention lack bite: it seems hard to take seriously the possibility that in the mirror world, the subject is perceptually attending to a mirror object. If so, then

---

<sup>8</sup> See Lewis [1999b, p. 64]. As Brian Weatherson pointed out, however, it is conceivable that a more substantial role for causation could be wrought out of Lewis' preferred formulation of charity (see fn. 10 below).

<sup>9</sup> Further, one could well imagine principled reasons to do with the mechanisms of perceptual attention for privileging some rather than other objects that are causally downstream from perceptual experience as the objects of attention.

we have leverage to eliminate crazy interpretations of the associated deictic thoughts and speech acts.<sup>10</sup>

Perceptual attention highlights the needs to give causal inputs a constitutive role in explaining why certain crazy interpretations are incorrect. It is not hard either to justify a constitutive role for causal outputs (actual or dispositional). For example, as Gareth Evans [1982. ch. 6] observed, it is highly plausible that causal relations to action will play a constitutive role in explaining how mental representations of space can be hooked onto places in actual space. Suppose, for example, that someone is lost in the dark and mentally represents a certain direction. What makes that representation – call it D -- a representation of a particular direction in real space? An adequate answer will arguably have to invoke the fact that when an intention to move ‘along D’ figures in the subject’s intentions, that will dispose her to move along a direction in real space. Absent this dispositional link, it is hard to see that ‘D’ could lock on to that real spatial direction at all.<sup>11</sup>

In short, any serious attempt to connect the crazy interpretation problem with foundational philosophy of mind will give causal constraints a much more central role than Lewis envisages.

---

<sup>10</sup> Of course there is still work to be done in fending off crazy interpretations of demonstratives that tie their referents not to the objects of attention but to something else (and hence which posit bizarre characters).

<sup>11</sup> There remain difficult issues connected to Kripke’s [1982] discussion of the content-fixing role of dispositions.

(2) In many cases, a crazy interpretation will do quite well on the score of charity, but less well on the score of knowledge. For example, on a natural version of a micro world crazy interpretation, the subject's dispositions to judge will not be sensitive to the goings on at the micro world and so his beliefs – though true --will be too 'lucky' to count as knowledge. Timothy Williamson [2004] has suggested that knowledge maximization should supplant charity at the foundations of semantics.<sup>12</sup> If we are willing to accord knowledge this constitutive role (or even to concede that *inter alia*, an interpretation should make people reliable belief-formers and not merely belief formers with a good *de facto* track record), then we can explain the inadequacy of a wide range of crazy interpretations. (Of course certain sufficiently well-crafted crazy interpretations will still remain live.)<sup>13</sup>

---

<sup>12</sup> I note in passing that Williamson is not after a set of principles that allow one to derive the intentional facts from a non-intentional groundfloor: '... present purposes do not require us to pursue the heroically ambitious quest for .... reduction. What we need are correct nontrivial principles about propositional attitudes that somehow link belief and truth, metaphysically rather than epistemologically. Such principles can fall far short of reducing the intentional to the non-intentional, even of fixing the supervenience of the former on the latter' (p. 138).

<sup>13</sup> Lewis' preferred formulation of charity is in fact in terms of 'reasonable belief'. If the relevant notion of reasonableness is glossed in subjective Bayesian terms, it is unclear whether this refinement on 'truth maximization' helps with the crazy interpretation problems I am concerned with. I am unsure as to what Lewis actually had in mind. For relevant discussion, see Lewis [1999a, p. 52].

(3) Given any context dependent expression, there will be a function from contexts to semantics values – call this its *character* – that represents how the content of that expression varies according to context.<sup>14</sup> Now when deciding between various candidate interpretations for a linguistic community, we will have to choose among various candidate functions that purport to represent how the expression’s content depends upon context. As Williams notes, the Lewisian might consider extending the eligibility idea to candidate characters: *inter alia*, interpretations that assign more natural characters to expressions are better. (Note that this constraint is not derivable from the original: there could be a candidate character that assigns highly eligible contents to a predicate on each occasion of use but where there was no natural rule for explaining how semantic content varies with context.)

(4) Lewis’ own picture is that there is a privileged set of perfectly natural properties that constitute a supervenience base for the world, and that other properties are graded by naturalness according to the length of definition by the perfectly natural properties.<sup>15</sup> I have complained elsewhere that the definitional conception of relative naturalness is unsuitable for Lewis’ purposes.<sup>16</sup> In the present context, it is worth underscoring just how inadequate that conception is to Lewis’ use of naturalness in semantics.

Let us take a representative case. One might hope that Lewis’ naturalness test, as applied to characters (see 3 above) could help to explain the correctness of the actual rule

---

<sup>14</sup> See Kaplan [1989]

<sup>15</sup> See, for example, Lewis [1999b, p. 66].

<sup>16</sup> See Hawthorne [2006].

for the pure indexical 'I' -- 'I' refers at a context to the speaker at the context. After all, that rule seems highly non-gerrymandered. But by Lewis's official lights, the rule is highly unnatural. We would, for example, require some definition of the property of being a speaker in the canonical language, one that delivers the set of speakers for any given possible world. But we cannot even do this in principle using just the language of microphysics, since the property of being a speaker has alien realizations (i.e. possible beings that are speakers but not made of the stuff of actual microphysics). It will be tempting to invoke causal and functional predicates in any first pass gloss at what it is for an object to be a speaker. But when one does not cheat and actually thinks through the 'canonical statement' of every causal predicate in the Lewisian system, one realizes that the canonical expression of every such predicate will be fantastically long -- after all, for Lewis, a Humean, facts about nomic and causal dependence are derived, holistic facts, whose canonical statement would be insurmountably complex. Remember also that the original statement of the rule for 'I' is vague (since terms like 'speaker' are vague.) Think now about what a precisification of the rule would look like. For every speaker at every possible world, there are a number of candidate overlapping objects that satisfy 'speaker' under some precisification, ones that differ slightly at their boundaries. It seems that a precisification of 'speaker' in a canonical language will have to, one at a time, select a candidate out from its competitors from every possible cluster of speaker-candidates. In sum, when one doesn't cheat, one quickly realizes that a precisification of the rule for 'I' cannot be achieved in the canonical language using a sentence of countable length. The intuitively simple rule for 'I' will thus come out as hopelessly gerrymandered when one takes seriously Lewis' preferred method for calibrating naturalness.

It appears, then, that the ‘more natural than’ relation cannot serve Lewis’ semantic purposes if it is tied to definitional length in the canonical language. We should thus be willing to give relative naturalness a life of its own, one that allows properties that are of equal definitional distance from the microphysical ground floor to be of radically unequal naturalness. (As Williams is aware, his Pythagorean thought experiment only makes direct trouble for eligibility proposals that are framed in terms of definitional length.)

(5) The charity-eligibility picture proceeds within a framework that expects there to be a graspable recipe for defining semantic properties in terms of the fundamental ones, or at least an a priori procedure that we can cotton on to for deriving the semantic ones from the fundamental ones. But perhaps the expectation that one can contrive such a general recipe or procedure is misconceived from the outset.

Certain threads in Lewis’ own work point towards such a perspective. Think of the declarative fragment of English as a set of ordered pairs of sentences and functions from contexts to truth values. Each sentence is a sentence of English; each companion function specifies the possible contexts at which the sentence is true. Lewis thinks that (ordinary vagueness aside), it is true that we use English rather than some competitor set of ordered pairs that associates different functions with the same sentences. Now it is quite clear that Lewis does not believe that there is any finitely specifiable account of what it is about us that determines us to be users of English. After all, he tells us that the contextual variability of truth values of sentences at contexts is fixed by features which include ‘standards of precision, the salience relations, the presuppositions....’ ([1998], p.29) and that these in turn are fixed by ‘the previous course of the conversation that is still going on at the context, the states of mind of the participants, and the conspicuous

aspects of their surroundings' ([1998], p. 29). He goes on to quote with approval Cresswell's remark that 'there is no way of specifying a finite list of contextual coordinates'. ([1998], p. 30) This all implies that by the lights of the definability rule, properties such as *truth-in-English* and *speaking English* will be fantastically gerrymandered properties. There is no prospect whatsoever of specifying a finite recipe for discerning whether a speaker is using English on the basis of the fundamental intrinsic and relational facts about that speaker.

(Note, then, that we should not have a general expectation that the character of an expression be particularly natural. There are some context-dependent expressions – paradigmatically, certain simple indexicals -- for which there is a simple graspable rule that describes how their semantic values depend on context. But for a large range of expressions the semantic value at a context is the product of myriad forces working in tandem, with correspondingly complicated character. We should thus not be excessively optimistic about how much can be pinned down by extending the eligibility desideratum to character.)

(6) Once we have relinquished the idea that relative naturalness is to be tied to ease of definability from the groundfloor, we are free to take a more elevated view of semantic properties and relations themselves. In particular, we should take seriously the idea that while semantic properties and relations don't occur at the groundfloor, they are very natural, not gerrymandered. After all, wouldn't it be very strange to concede that such properties and relations as *is true*, *semantic expresses* and so on, were all massively gruesome? Once one embraces the naturalness of semantic properties, the observations under (5) above will not appear threatening. Granted, there no recipe for generating the

semantic properties from the fundamental ones. But given that naturalness does not have to be tied to ease of definability, that does not indict the naturalness and importance of semantic properties.<sup>17</sup>

\$4 One theme that Williams emphasises is the central place of the interpreter in Lewis' account of the foundations of semantics. According to Williams, this is driven by an 'interpretationalist' metasemantics according to which, inter alia,

Overall... semantic facts emerge because they are part of a simple, finitary theory whose predictions mesh with facts about utterance-conditions for sentences'.

....there need be no reference relation, characterisable in non-semantic terms, onto which our semantic vocabulary latches'

.....we explain the constitution of semantic facts, such as 'Londres' refers to London by appeal to holistic properties of a theory in which this claim figures.<sup>18</sup>

---

<sup>17</sup> Of course, there may still be true and interesting principles in the vicinity of charity and eligibility, even though those principles do not suffice to explain the naturalness of the semantic and even though those principles do not suffice to provide a recipe for deriving the semantic from the groundfloor (cf. fn. 12).

<sup>18</sup> Williams, this journal issue, p.

It is not clear to me how to build a plausible and distinctively ‘interpretational’ approach to metasemantics out of these remarks. First, since it makes no sense to say that there is no reference relation, the key phrase in the second remark is ‘characterizable in non-semantic terms.’ But it is hardly a distinctive thesis of interpreter-centric semantics that there is no finite way of specifying the intension of ‘refers’ and other semantic vocabulary in a fundamental language appropriate to the metaphysical groundfloor. Second, banal counterexamples to behaviourism will refute an interpretationalism that claims that utterance conditions for sentences provide a supervenience base for semantic content.<sup>19</sup> Third, there is at least a danger of regress in appealing to theories in which claims about reference figure since that requires saying what it is about a sentence in a theory that makes it a claim about reference. Fourth, so-called interpretationalist proposals can generally be rearticulated in a way that eliminates reference to the interpreter altogether, in the form of holistic constraints on how semantic facts supervene on non-semantic ones.<sup>20</sup>

---

<sup>19</sup> That emphasis also ignores the question of what constitutes hearer’s understanding as opposed to speaker’s meaning.

<sup>20</sup> In personal communication concerning a draft of this paper, Williams indicated that interpretationalism, as he understands it, is consistent with the thesis (described below) that the interpreter has a purely picturesque role in the metaphysics of semantic properties. I do not wish to speculate further about how exactly his remarks on interpretationalism were intended to be understood. Obviously, I may have taken some of them in the wrong spirit.

As a prelude to assessing the importance of the interpreter to metasemantics, let me offer my own (admittedly incomplete) taxonomy of the possible roles for an interpreter the foundations of semantics.

First, the role of the interpreter might be purely picturesque, where the question ‘How would an interpreter figure out the semantic facts from the non-semantic facts?’ is just a stylistic variant on the question ‘How do the non-semantic facts fix the semantic facts?’. This is the role of the radical interpreter in ‘Radical Interpretation’:

By what constraints, and to what extent, does the totality of physical facts about Karl determine what he believes, desires and means? To speak of a mighty knower, who uses his knowledge of these constraints to advance from omniscience about the physical facts P to omniscience about the other facts determined thereby, is a way of dramatizing our problem – safe enough, so long as we can take or leave it alone. The real life knower has all the problems of our fictitious knower, and more besides since he does not have all of P to draw on..... But these further obstacles to his investigations are irrelevant to our real topic (Lewis, 1983, p.111).

Lewis seems to be at pains here to distance himself from interpreter-centric metasemantics. Is there any reason to think Lewis had a serious change of heart? I don’t see it – the eligibility constraint is naturally understood as a proposed addition to the constraints of which Lewis speaks in ‘Radical Interpretation’.

A second reason why one might emphasise the interpreter is by way of a (brutish) insistence that the facts that are readily open to view by ordinary interpreters must determine the propositional attitudes and meanings of a speaker (and hearer).

A third reason for emphasising the interpreter would be provided by a commitment to the thesis that if semantic facts supervene on fundamental facts then they must supervene a priori on the fundamental facts. Such a view, developed inter alia by David Chalmers [1996] and Frank Jackson [1998], requires that the semantic facts cannot outrun those semantic conclusions that an idealized version of ourselves could infer from a complete statement of the physical groundfloor. The constraints and limits on the semantic facts are thus constitutively tied to the capacity of an idealized interpreter to recognize them. This motivation, like the second, is driven by neo-verificationist commitments.

A fourth, and to my mind most interesting reason for emphasising the interpreter, requires no such commitment. Let us grant that there are facts about which sentences are true at which contexts, which properties predicates express at which contexts, and which objects singular terms refer to at which contexts. Even so, there remains an interesting philosophical challenge concerning the significance of those semantic properties and relations. There are, after all, a variety of properties and relations  $x$  *true\** at  $y$ ,  $x$  is *true\*\** at  $y$ ,  $x$  *expresses\**  $F$  at  $y$ ,  $x$  *expresses\*\**  $F$  at  $y$ ,  $x$  *refers\** to  $y$  at  $z$ ,  $x$  *refers\*\** to  $y$  at  $z$ , and so on, that differ only slightly in their pattern of instantiation from *true at*, *expresses at* and *refers at*. That is, most of the expression-context pairs that instantiate the *refers at* relation also instantiate the *refers\* at* relation; and so on. The challenge proceeds by inquiring as to why *truth at*, *expresses at* and *refers at* are important, when their starred

competitors are not?<sup>21</sup> One might even try to extend the interrogation to certain properties of propositions. Granting the monadic properties of truth and falsity of propositions, what makes *those* properties more important than truth\* and falsity\*? (Obviously, one can raise analogous questions about validity and validity\*.)

Let us distinguish two styles of answer to the inquiry.

Examples of style one (of varying levels of plausibility): (a) *Truth at* (and, relatedly, *truth* of utterances) is much more natural than *truth\* at* (and *truth\** of utterances). Further, the monadic properties of propositional *truth* and *falsity* are much more natural than their counterparts *true\** and *false\**. (b) Truth figures in bona fide causal explanations but truth\* doesn't. (c) It is a basic fact about the world that knowledge (which requires truth) is intrinsically valuable, while knowledge\* (which requires truth\*) is not.

Examples of style two. (a) There are all sorts of questions I am interested in – Whether it will rain tomorrow, whether I am happy, etc. Now I know that I am happy iff 'I am happy' is true at this context. And I know that it will rain tomorrow iff 'It will rain tomorrow' is true at this context. So I am, obviously, interested in which sentences are true at my context. But the following equivalence does not hold with full generality: *s* is true at a context iff *s* is true\* at a context. So I cannot justify an interest in truth\* in that way. (b) Take any proposition *P* such that I want to know whether *P*. I know that *P* iff It is true that *P*. So, obviously, it is also going to be of interest whether it is true that *P*. By hypothesis, that equivalence does not hold with full generality for true\*. So I cannot justify an interest in truth\* in that way.

---

<sup>21</sup> Cf.. Stich [1990].

The details of the responses are not especially crucial. (There are clearly holes that need to be filled.) The pertinent difference between the two styles of responses is that the latter class *uses* various semantic and intentional descriptions of ourselves in order to justify its claims about the significance of truth over truth\* and so on, whereas the former class does not. If a style two response is the best or most appropriate, then it begins to seem that the only way to articulate the significance of semantic and intentional properties is by making use of our semantic and intentional self-conceptions. If so, there is something deeply right about the idea that the significance of semantic properties can only be appreciated by someone who already has the facility for self-ascribing those properties using the machinery of our ordinary semantic concepts.

Suppose there were a being who was omniscient about the metaphysical groundfloor and who had labels for all the supervenient relations, but who lacked concepts that at once represented the semantic properties while also playing the conceptual roles of our own semantic concepts, and also lacked concepts that at once represented various propositional attitudes while playing something like the conceptual roles of our own propositional attitude concepts. Those attracted to style two answers would expect such a being to have nothing useful to say about what makes semantic properties more significant than properties and relations in the vicinity. Our omniscient being will label the semantic properties and relations all right, but will have nothing to say as regards their significance. By contrast those who hope for a style one response aim for an answer to our challenge that can, as it were, be articulated from the ‘view from nowhere’.

Obviously there is a great deal more to be said. But I hope to have done enough to justify my suspicion that one reason – and potentially a laudable one – for keeping the interpreter in view in metasemantics is that the significance of semantic properties may be something that cannot be accounted for in straightforwardly metaphysical terms.

## REFERENCES

- Chalmers, D. (1996). *The Conscious Mind*. Oxford University Press.
- Evans, G. (1982). *Varieties of Reference*. Oxford University Press.
- Hawthorne, J. (2006). 'Quantity in Lewisian Metaphysics,' In *Metaphysical Essays*. Oxford University Press, p. 229-237.
- Jackson, F. (1998). *From Metaphysics to Ethics: A Defense of Conceptual Analysis*. Oxford University Press.
- Kaplan, D. (1989) 'Demonstratives,'. In Almog, Perry and Wettstein, eds., *Themes from Kaplan*, Oxford University Press, p. 481-566.
- Kripke, S. (1982) *Wittgenstein on Rules and Private Language*. Harvard University Press.
- Lewis, D. (1969). *Convention*. Harvard University Press.
- Lewis, D. (1983). 'Radical Interpretation.' Reprinted in *Philosophical Papers Volume 1*, Oxford University Press, p. 108-122. First published in 1974.
- Lewis, D. (1998). 'Index, context, and content.' In *Papers in Philosophical Logic*, Cambridge University Press, p. 21-44. First published in 1980.

Lewis, D. (1999a). 'New Work For a Theory of Universals.' Reprinted in *Papers in Metaphysics and Epistemology*, Cambridge University Press, p.8-55. First published in 1983.

Lewis, D. (1999b). 'Putnam's Paradox'. Reprinted in *Papers in Metaphysics and Epistemology*, Cambridge University Press, p.56-77. First published in 1984.

Quine. W. V. (1960). *Word and Object*. MIT Press.

Stich, S. (1990). *The Fragmentation of Reason*. MIT Press.

Williamson, T. (2004). 'Philosophical 'Intuitions' and Scepticism about Judgment'. In *Dialectica* 58, p. 109-153.