

Modal Rationalism & Logical Empiricism: Some Similarities

Stephen Yablo
MIT

I. Introduction

You remember logical empiricism, otherwise known as logical positivism. It was the big thing for a while in the middle of this century. A number of ideas were involved: the verifiability criterion of meaningfulness; the analytic theory of the a priori, the idea that a sentence's meaning was a matter of the experiences counting for and against it; and so on. As the record clearly shows, most of these ideas came eventually to seem mistaken. And mistaken for the same sort of reason. The problem with positivism is that it ran semantics and epistemology too closely together.

So, for example, why should meaningfulness, a semantic notion, depend so heavily on verifiability, an epistemological one? Why should a priority, an epistemological notion, depend so heavily on analyticity, a semantical one? Not only are the envisaged connections puzzling in principle, but there look to be counterexamples in both cases. "God exists" strikes many people as meaningful even if unverifiable. " $2=2=4$ " strikes many of us as apriori though not analytic.

All that is ancient history, however. Why care about these old ideas today? Because, or at least this is my conjecture, the old ideas are still around. A theory of meaning advocated by David Chalmers, Frank Jackson, and David Lewis -- and used by the first now and the second until recently as a basis for rejecting physicalism -- is caught up in something very like the old positivist themes. It's caught up in particular with the reductionist idea that what a sentence means goes hand in hand with the lower-level circumstances that would verify it.

Now this may strike you as the farthest thing from positivism. It may strike you as just the truth-conditional theory of meaning, which is widely viewed nowadays as just common sense, and objectionable only for being so toothless. If you are right about that, then my objection is to the truth-conditional theory of meaning as well. I'll be saying that theory too has more in common with logical empiricism than is usually acknowledged.

II. Carnap on Internal/External

The aspect of logical empiricism I want to focus on is best known from Carnap's ESO, and Quine's criticisms of ESO in "Two Dogmas." The ostensible topic here is Carnap's distinction between internal existence questions like "are there primes between 10 and 20?" and external existence questions like "are there numbers?" But as we'll see, the issue quickly turns to reductionism.

Here's what Carnap says. Meaningful talk about Xs has got to proceed under the auspices of a linguistic framework, that is, a set of "rules for forming statements [about Xs] and for testing, accepting, or rejecting them." An ontologist who respects this requirement by querying "the existence of [Xs] within the framework" is said by Carnap to be raising an internal existence-question. □ But it is not only existence-questions and –statements that can be classified in this way. The idea is meant to be perfectly general. Any question raised within a framework, and answerable to that framework's rules, is in Carnap's sense internal.

The point Carnap wants to make about internal questions -- 'is it true that S?' -- is that they raise no difficulties of principle. It is just a matter of whether applicable rules authorize you to say that S. If they do, the answer to 'is it true that S?' is yes; otherwise no; end of story. □

Once again, the application that interests Carnap is to internal existence-questions. Looking at what he says about these will help us to understand more about how frameworks work. 'Do numbers exist?' understood internally is easily answered: of course they do, this follows from the fact that there are prime numbers between 10 and 20.

That the answer comes so easily, however, shows that the internal existence-question is not the one the philosopher meant to be asking: it is not the "question of realism." A system of rules making "there are material objects" or "there are numbers" unproblematically assertible is a system of rules in need of external validation, or the opposite. Are the rules right to counsel acceptance of "there are Xs"? It is no good consulting the framework for the answer; we know what it says. No, the existence of Xs will have to be queried from a position outside the X-framework. The philosopher's question is an external question.

Now, Carnap respects the ambition to cast judgment on the framework from without. He just thinks philosophers have a wrong idea of what is coherently possible here. How can an external deployment of a statement about Xs mean anything, when by definition it floats free of the rules whence alone meaning comes?

There are meaningful questions in the vicinity. But these are questions that mention "X" rather than using it: e.g., the practical question "should we adopt a framework requiring us to use "X" like so?" □ If the philosopher protests that she meant to be asking about Xs, not the term "X," Carnap has a ready reply: "You also thought to be asking a meaningful question, and one external to the X-framework. And it turns out that these conditions cannot be reconciled.

III. Quine's charge of reductionism

Now let's look at Quine's response to Carnap. The key to Carnap's position (as he sees it) is that "the statements commonly thought of as ontological are proper matters of contention only in the form of linguistic proposals." □ But now, Quine reasons, similar claims have been made about the statements commonly thought of as analytic; theoretical-sounding disputes about whether, say, the square root of -1 is a number are best understood as practical disputes about how to use "number." So, idea: the external existence-claims can be (re)conceived as the analytic ones. The objection thus looks to be one of guilt-by-association-with-the-first-dogma: "if there is no proper distinction between analytic and synthetic, then no basis at all remains for the contrast which Carnap urges between ontological statements and empirical statements of existence." □

Unfortunately, the association thus elaborated doesn't look all that close. For one thing, analytic existence-claims show no particular tendency to be external. Quine appreciates this but pronounces himself unbothered:

there is in these terms no contrast between analytic statements of an ontological kind and other analytic statements of existence such as "There are prime numbers above a hundred"; but I don't see why he should care about this.

Quine's proposal also deviates from Carnap in the opposite way; existence-claims can fail to be analytic without (on that account) failing to be external. An example that Carnap himself might give is "there are material objects." Quine apparently considers it a foregone conclusion that we will have experiences making "there are material objects" assertible in the thing framework. □ But how could it be? It is not analytic that experience even occurs. □

So far, so confusing. To add to the mystery, Carnap appears to agree with Quine that the distinctions are linked – that internal/external depends on analytic/synthetic:

Quine does not acknowledge [my internal/external] distinction because according to him "there are no sharp boundary lines between logical and factual truth, questions of meaning and questions of fact, between acceptance of a language structure and the acceptance of an assertion formulated in the language."

Since internality is just a different thing from analyticity, the link if any has got to take a different form than we have yet seen. There is a clue in the parallel Quine draws between, on the one hand, meaning vs. fact, and on the other hand, acceptance of a language structure vs. acceptance of an assertion formulated in the language.

Accepting a language structure, Quine and Carnap agree, is accepting a bunch of meaning-giving assertion rules. These rules are not statements, and so there is no question of calling them analytically true. But they are the nearest thing to, namely, analytically valid or correct. The rules are what give X-sentences their meanings, hence they "cannot be wrong" as long as those meanings hold fixed. The upshot is that analytic/synthetic does define internal/external: not directly, by providing an outright equivalent, but indirectly through its role in the notion of a framework.

Now we have the sought-after link: internal/external presupposes analytic/synthetic by presupposing frameworkhood; for frameworks are made up inter alia of analytic assertion rules. Someone might ask, "why should analytic rules be as objectionable as analytic truths?" But that is essentially to ask why Quine's second dogma -- the reductionism that finds every statement to be linkable by fixed correspondence rules to a determinate range of confirming observations -- should be as objectionable to him as the first.

The objection is the same in both cases. Any observation can work for or against any statement in the right context. Hence no assertion or rule of assertion can lay claim to being indefeasibly correct, as it would have to be were it correct as a matter of meaning. Quine may be right that the two dogmas are at bottom one; still, the charge narrowly drawn is one of guilt-by-association-with-the-second-dogma, the dogma of reductionism.

IV. Reductionism and defeasibility

Now let's move on to a further irony of the debate. Quine has just objected to Carnap's meaning rules that no assertion or rule of assertion is indefeasibly correct. The irony is that Carnap agrees that no assertion or rule of assertion is indefeasibly correct. He says already in the 1930s that "all rules are laid down with the reservation that they may be altered as soon as it seems expedient to do so" (Carnap 1937, 318). Carnap agrees with Quine that any link between theory and observation can be broken, and any can in the right context be forged.¹

The thing is that he puts a different spin on these scenarios. There is indeed (thinks Carnap) a possibility that can never be foreclosed. But it is not the possibility of our correcting the rules to accommodate some new finding about the conditions under which X-statements are "really true";² it is that we should decide for practical reasons to trade the going framework for another, thereby imbuing "X" with a new and different meaning.³

So while Quine and Carnap agree that the rules are up for grabs, Carnap thinks that the rules are meaning-giving; to change them is not to revise our doctrine but only to switch to an alternative (if similar-sounding) vocabulary. To the extent that Carnap limits this point to artificial languages and frameworks, Quine can agree. But when it comes to languages in use, Quine sees no basis for Carnap's meaning/doctrine distinction; and so he maintains that virtually any change in the rules can and should be seen as having a doctrinal element. It's not that we give "X" a different meaning, it's that we change our minds about where the Xs are.

V. Modal rationalism

That's enough for now about the Carnap/Quine debate. Briefly, Carnap thinks it's a semantical matter under which observational conditions it's to be asserted that S, whereas Quine thinks it's a doctrinal matter. It's Carnap's position that I say bears comparison with modal rationalism.

Here to a first approximation is the modal rationalist picture. Associated with every sentence S are two propositions or sets of worlds. The first, the primary intension, is

|S|₁ = the set of all worlds 'making S true'
= the set of all w such that if w is actual then S is true.

This gives the a priori aspect of meaning; one knows just in virtue of understanding S what the actual world has to be like for S to be true. □ The secondary intension is

|S|₂ = the set of worlds 'satisfying' S
= the set of all w that S truly describes.

This gives the 'what is said' aspect of S's meaning. To the extent that what is said depends on contextual factors, one doesn't know S's secondary intension just in virtue of understanding S.

1
2
3

Suppose for instance that S is "water is plentiful." Then what S in fact says, assuming that the watery appearances around here are appearances of H₂O, is that H₂O is plentiful. The secondary intension is thus the set of worlds with lots of H₂O. But what S requires of this world is different; it requires that the watery stuff be plentiful. So the primary intension is the set of all worlds with lots of watery stuff. And now here are the fundamental 'predictive' theses:

(A) it is a priori that S iff $|S|_1$ is the set of all worlds, that is, iff S is true in virtue of the a priori aspect of its meaning.

(N) it is necessary that S iff $|S|_2$ is the set of all worlds, that is, iff what S says is the case must be the case.

(A) and (N) neatly explain how S can be necessary yet aposteriori, or contingent yet apriori. But the point that I want to focus on comes earlier, in the intuitive explanation of primary intension as the a priori aspect of meaning. To call primary intension the apriori aspect of meaning is to say in effect that

(P) whether a world w verifies S is determined by the apriori aspect of its meaning.

The word 'verified' here means just: is such that S is true if w in fact obtains. No connection is intended with the logical empiricist doctrines of verificationism and reductionism. Or rather, no connection is imputed. I suspect that there is a connection as I now proceed to explain.

VI. Initial comparisons

First let's clear away some apparent differences between modal rationalism and Carnap-style reductionism. Reductionism envisages analytic rules connecting theory to experience. MR doesn't envisage anything like that. Yes, people have to be told a priori whether S is true in a presented world. But Chalmers and Jackson don't say that the world is presented experientially. There's no question then of associating MR with an experiential or phenomenalist reductionism.

This is fair enough. But if one looks at Carnap's writings on "protocol sentences," it turns out that LE reductionism does not have to be terribly experiential either. Under the influence of Neurath, Carnap thinks it is somewhat of an open question which sentences ought to be counted as protocols. Sometimes a protocol-sentence is said to be any sentence "belonging to the physicalistic system-language" which we are prepared to accept without further tests (Ayer, 237). Usually it is said to be a matter of convention which sentences will count as protocols. The important point for our purposes is that Carnap thinks there are analytic rules connecting theoretical statements with protocols whatever protocol turn out to be.

Another seeming difference between (logical) empiricism and (modal) rationalism emerges from Quine's complaint that reductionism overlooks the holistic nature of confirmation. The complaint might be this: One never knows whether S is really correct until all the observational evidence is in, which it never is. Hence any rules portraying S as verifiable on the basis of limited courses of experience -- courses of experience small enough to be enjoyable by particular observers -- would be untrue to the way confirmation actually works. This complaint the (modal) rationalist can rightly claim to have avoided. The (m) rationalist never represents limited information about a world as enough to ensure that S; the rules he contemplates take as input complete information. Dave is very clear about this:

[Quine says that] purported conceptual truths are always subject to revision in the face of sufficient empirical evidence. For instance, if evidence forces us to revise various background statements in a theory it is possible that a statement that once appeared to be conceptually true might turn out to be false.

This is so for many purported conceptual truths, but it does not apply to the supervenience conditionals that we are considering, which have the form "If the low-level facts turn out like this, then the high-level facts will be like that." The facts specified in the antecedent of this conditional effectively include all relevant

empirical factors-.The very comprehensiveness of the antecedent ensures that empirical evidence is irrelevant to the conditional's truth-value (TCM, 55).

But there is another way of construing Quine's complaint that makes it quite unrelated to the "total evidence" issue. Quine says that

the dogma of reductionism survives in the supposition that each statement, taken in isolation from its fellows, can admit of confirmation or infirmation at all. My countersuggestion, issuing essentially from Carnap's doctrine of the physical world in the Aufbau, is that our statements about the external world face the tribunal of sense experience not individually but only as a corporate body (TDE, 41).

Never mind the irony of Quine attacking Carnap for overlooking a point he himself gets from Carnap. Quine's complaint here is not that S's confirmational status is underdetermined until all the empirical evidence is in; it's that S's confirmational status is not fully determined even by the full corpus E of empirical evidence. The degree to which E confirms S, Quine thinks, is tied up with the extent to which E or aspects of E are deducible from S. But nothing of an observational nature is deducible from S except with the help of a background theory T. Hence the degree of support that E lends to S depends on which background theory we use.

This complaint would be easily evadable if there were an analytically guaranteed fact of the matter about which theory T is selected for by E. For in that case one could simply ask whether E supports S relative to the E-preferred theory whatever it might be.

One has to assume, then, that this is what Quine is really concerned to deny; he denies that there are analytic connections between the total corpi (?) E of empirical evidence and theories T of nature, and hence there are no analytic connections between E and particular statements S. A number of things suggest that these analytic confirmational relations are indeed the target:

I am impressed, apart from prefabricated examples of black and white balls in an urn, with how baffling the problem has always been of arriving at any explicit theory of the empirical confirmation of a synthetic statement (TDE, 49)

This could be taken to mean just that the sought-after theory of confirmation would have to be very very complicated. But Quine has something different in mind. He is aware after all of Carnap's attempts to work out a logic of confirmation which would tell us what to believe on the basis of such and such observations. He is aware too that the attempt failed even for the simplest sort of examples. Carnap came up with a bunch of confirmation functions (the m-functions) and was forced to regard the choice between them, and hence between the statements they recommend, as one that could not be made on logical or semantical grounds.

What then is Quine's point? His point is underdetermination: "total science, mathematical and natural and human, is underdetermined by experience" (TDE, 45). The version of underdetermination he needs is really a rather mild one. He needn't deny that there is an objectively best theory relative to a given body of evidence. He needn't even deny that there's a single most rational theory to adopt. All he need claim is that the choice between theories compatible with the evidence cannot be made on purely analytic grounds. Our choice of theory

turns on our vaguely pragmatic inclination to adjust one strand of the fabric of science rather than another. Conservatism figures in such choices, and so does the quest for simplicity.

This can be reconciled with the analytic view of confirmation relations only by supposing that meanings suffice to determine for us how conservative we should be, and what sort of simplicity to favor, and how to trade these sorts of desiderata off against each other. The idea would have to be that when people differed on what sorts of tradeoffs made sense, at most one could be considered to be speaking correctly, i.e., in accordance with the meanings of her words. But that, it seems clear, is not how the science game is played.

VII. How we operate

This brings us to one final irony. Carnap agrees that that's not how the science game is played. His goal as he usually describes it is not to characterize the true nature of meaning, but to give us tools for making our discursive

practice more rational and more efficient. He thinks disputants should pick a common framework and then resolve their disagreements by reference to its assertion rules. Here is how he puts it in Testability and Meaning:

It seems to me that it is preferable to formulate the principle of empiricism not in the form of an assertion ... but rather in the form of a proposal or requirement. As empiricists we require the language of science to be restricted in a certain way... (TM, #27)

Based on passages like this, one recent commentator has summarized Carnap's view as follows:

Criticisms of the meaning/belief distinction rest on the lack of a principled criterion for [semanticity] – that no empirical method can be found for making it. However, for Carnap, such a distinction is to be reached by agreement in a conflict situation. Maximize agreement on framework issues and situate disagreement on either empirically answerable problems or on questions of a pragmatic nature about the framework (CTDE, 1026)

One can argue about whether this would really be helpful. All I'm saying right now is that not even Carnap believes that it's how we really operate: that our actual practice lends itself to a distinction between semantic factors in assertion and doxastic ones.

Is there anyone who does believe that it's how we operate? It seems to me that the modal rationalist believes it, or at least that such a view is implied by modal rationalism. Here is a first stab at convincing you of this. The modal rationalist says that grasp of S's meaning, or at least the kind of grasp you need to count as understanding S, is knowing which worlds w are such that if this is w, then S. This applies not just to observation-level statements, but theoretical statements T as well. It would seem to follow that the rationalist has to see it as part and parcel of knowing T's meaning to know what the world has to be like for it to be the case that S. And that is not obviously different from Carnap's idea of analytic confirmation rules.

I say "not obviously different" because it's unclear to begin what form this "knowing which worlds are S-worlds" is supposed to take. Under what description exactly are we supposed to think of the relevant set of worlds?

One idea about this is clearly not intended. It's not supposed to be that my grasp of the meaning of "there are horses" is lies in the fact that I know it to be true in the worlds with horses, or rather the ones such that if they are actual then there are horses. For suppose that a 'homophonic' grasp of the set of verifying worlds was all one needed. Then contrary to (A) above, there would be no reason at all to expect a sentence to be knowable a priori just because its primary intension contained all worlds.

The easiest way to explain this is with an example from Chalmers's discussion of physicalism. Consider the conditional "if things are physically like so, then there is pain." According to (N), the only way for this to be non-apriori is for there to be worlds not in its primary intension: there have got to be zombie-worlds. But if a homophonic grasp of the primary intension is all that's required, then we can explain the non-apriority of the conditional without bringing in zombie-worlds at all.

It'd be enough to point out that to understand the conditional it suffices to understand its antecedent A and consequent C, ie., to know that their primary intensions are the A-worlds and C-worlds respectively. And to know that much about the primary intensions you don't need to know the set-theoretic relations between them. Even if the physically-like-so worlds were included in the pain-worlds, there'd no reason why the understander should realize this, and hence no reason to think she should be able to determine the conditional's truth value just on the basis of her understanding.

What the modal rationalist intends then is that to understand S one should know what lower-level facts have to obtain in a world if it is to verify S. And this is indeed what Chalmers says:

The supervenience conditionals that we are considering, which have the form "If the low-level facts turn out like this, then the high-level facts will be like that." The facts specified in the antecedent of this conditional effectively include all relevant empirical factors...

These conditionals are for the modal rationalist analytic in the sense of being true in virtue of meaning; indeed they are true in virtue of the aspect of meaning to which we have a priori access. This is why I say that the modal rationalist is committed to something *prima facie* analogous to the analytic confirmation relations advocated by Carnap and rejected by Quine. The modal rationalist who wants to escape Quine's criticisms has got to do one of two things:

(a) show that the criticisms don't work even against logical empiricism;

(b) show that the cases are relevantly different.

To accomplish (a) would be to find a mistake in Quine's reasoning. Maybe, for example, it's just untrue that theory is underdetermined. To accomplish (b) would be to show that what the modal rationalist says is different enough from what the logical empiricist says that the Quinean critique doesn't generalize. Maybe, for example, the lower-level facts on the basis of which we can tell a priori whether S are quite unlike the "empirical" facts on the basis of which we can't tell a priori whether S. The rest of the paper explores in a roundabout way the prospects for doing either of these things, and thereby escaping Quine's critique.

VIII. Reference-fixing

I've been saying, basically, that the rationalist's analytic supervenience conditionals are hard to square with the Quinean doctrine of 'underdetermination of theory.' The version of underdetermination we need is a mild one. All we need claim is that the choice between theories compatible with the evidence cannot be made on purely analytic grounds: considerations about which theory is simpler or clearer or more natural inevitably come into play. This alone makes it hard to see how the supervenience conditionals can be analytic. Because the issue of which supervenience conditionals are correct cannot be separated from the issue of which theory is correct given the facts in the conditional's antecedent. If the one issue is not analytically resolvable, then neither is the other.

One way that the rationalist tries to weaken our resistance to the analyticity claim is by bringing in the notion of a reference-fixing condition or description. Suppose for argument's sake that water is stipulated to be the local clear drinkable stuff. (This is an idealization but it is said a harmless one.) Then just by virtue of understanding 'water,' I know what its extension is relative to any assumption about the lower-level facts.

I wonder though whether the idealization is really so harmless. To call it harmless suggests that even an unidealized understanding of 'water' would determine the word's extension relative to any hypothesis about the lower-level facts -- so that if someone disagreed with me about where the water was in a world specified in lower-level terms, that would show she meant something different by water. Is this plausible?

It seems to be agreed all around that the extension of 'water' in a world is at least somewhat a function of what the right scientific theory is. But then if meanings aren't enough to single out the right scientific theory, they are unlikely to single out the right extension for 'water' either. It would seem to follow that the reference-fixer for 'water,' to the extent that it's true in virtue of (an a priori accessible aspect of meaning), cannot possibly give explicit necessary and sufficient conditions in lower-level terms; there just aren't any lower-level conditions that are analytically guaranteed to deal correctly with all scenarios, no matter how outre. A reference-fixer guaranteed to give the right results cannot avoid appealing at some point to notions like naturalness, simplicity, nonarbitrariness, etc. -- in a word reasonableness. But then the supervenience conditionals are not analytic unless it's analytic that such and such is the most reasonable thing to say about an imagined scenario.

Now there are places where Dave seems to agree that considerations of reasonableness will have to come in:

In certain cases, the decision about what a concept refers to in the actual world involves a large amount of reflection about what is the most reasonable thing to say; as, for example, with questions about the reference of 'mass' when the actual world turned out to be one in which general relativity is true, or perhaps with questions about what qualifies as 'belief' in the actual world-consideration of just what the primary intension picks out in various actual-world candidates may involve a corresponding amount of reflection. But this is not to say that the matter is not a priori: we have the ability to engage in this reasoning independently of how the world turns out. (58).

This sounds right; we do have some ability to engage in this reasoning independently of how the world turns out. The question is whether the results of our reasoning deserve to be counted into the meanings of words like 'mass' and 'belief.'

Of course one is free to define 'meaning,' or 'apriori component of meaning,' however one wants. But we need to be clear that this is somewhat of a novel use. The usual view is that differences about the application of a word in specified circumstances sometimes reflect differences in meaning, sometimes due to differences in doctrine: differences about what it's reasonable on balance to believe. The effect of modal rationalism is to eliminate this second category, so that people who have different ideas about what (it's reasonable to think) various people in fact believe in a given situation come out meaning different things by the word 'belief.' This makes an enormous difference to the significance of the rationalist's claim that the supervenience conditionals are analytic or true in virtue of meaning.

IX. The meaning of 'reasonable'

But we are getting ahead of ourselves. The question we need to be asking is, how far should differences about when to apply a word be taken to reflect differences in meaning? It helps to focus on the word 'reasonable' since it is conflicting judgments about reasonableness that underlie disagreements about the extensions of other words.

So: is it the case that anyone who means what I do by 'reasonable' has got to agree with me on the reasonableness or not of proposed courses of action? Should I regard it as an analytic consequence of the lower-level descriptive facts that this course of action is the proper one to pursue?

I can't regard it as an analytic consequence, it would seem, without falling afoul of the open question argument. Imagine someone who says, 'that this alternative is descriptively so and so entails that it is the alternative that makes the most rational sense.' This person can be accused of not getting it. It flows from the concept of reasonableness that it's always in order to wonder whether the proposed option is really the most reasonable one. The question is, as they say, always open. And an open question, whatever else may be true of it, is a question whose correct answer is not analytically true.

So: the fact that one can't avoid appeals to reasonableness in deciding extensions shows that the the supervenience conditionals are not analytic. Or rather, it shows this if we agree that the open question argument prevents there from being an analytically determined fact of the matter about which option is rationally warranted.

X. The open question argument

Given our reliance here on the open question argument, it has to be counted as a marvelous stroke of luck that at least one modal rationalist, Frank Jackson, has considered this argument and rejected it. Jackson is responding to the argument as it concerns 'right' rather than 'reasonable.' He thinks that it's not an open question what is the right thing to do in a given set of descriptive circumstances. Which actions are right is analytically determined by the lower-level descriptive facts. First let's see how it's determined; then we can ask how the open question is supposed to be closed.

Take just about any predicate F, Jackson thinks, and the following is analytic: Fness is the property playing the F-role in folk F-theory. This holds in particular for the property of rightness -- with the twist that it is not current folk morality that calls the shots but a "maturer" version. Reference is pegged to the place "where folk morality will [would] end up after it has been exposed to debate and critical reflection" (133). It is analytic he thinks that an action is right iff it has whatever property D meets descriptive condition M, a condition extracted by the Ramsey-Lewis method from mature folk morality.

This would seem to make Jackson easy prey for the open question argument: "I see that keeping my promise has the M-property, but is it right?" He has two main replies.

(1) Since "what matters is the nature of mature folk morality, there will, here and now, inevitably be a substantial degree of 'openness' induced by the very fact that the rightness role is currently under investigation" (151).

(2) If the question still seems open when "all the negotiation is over and we have arrived at mature folk morality, we-are entitled to dig in our heels and insist that the idea that what fits the bill that well might fail to be rightness is nothing more than a hangover from the platonist conception that the meaning of 'right' is somehow a matter of its being mysteriously attached to the form of the right" (152).

Suppose we consider these in turn. If we accept reply (1), then we ground the openness of "is that right?" in the openness (for us today) of "is that mature folk morality?" Jackson talks as though this latter openness were a matter of ordinary descriptive ignorance about the future; we haven't got there yet, so we don't know.

But of course morality could take a direction that we would regard as quite misguided, and for good and specific reasons; tomorrow's moralists might end up resting a lot on tendentious analogies, or giving in to self-interest, or etc. The point is just that it is no part of current folk morality to defer to whatever comes along. If deference must be paid, it is owed to those (probably hypothetical) future populations who had thought things through carefully, reasonably, and with due concern for all. So the best (1) can hope to accomplish is to ground the openness of "is it right" in that of "is that the most careful, reasonable, etc. way of developing the theory from here?"

But "careful," "reasonable," 'due concern' and so on are themselves evaluative terms. So Jackson faces a choice. Either he maintains about reasonableness too -- let it stand in for the rest -- that it is analyzable in descriptive terms, or he treats it as irreducibly evaluative.

One assumes that he will stick with his earlier descriptivism. "Reasonable" picks out whatever property plays the reasonableness-role in mature folk reason-theory, the theory we would arrive at if existing folk reason-theory were developed along the most reasonable lines. But this is circular; "the most reasonable lines" are the lines that would be identified as such by the theory lying at their terminus.□ So it is hard to see reply (1) as helping matters; if anything it confirms the impression that the meanings of 'right' and 'reasonable' leave their extensions a matter for debate.

Of course it's agreed that their extensions are a matter for debate now. But Jackson thinks that our current moral practice is predicated on the idea that the debate continued along present lines will eventually reach a conclusion. He says it is "part of current folk morality that convergence will or would occur. We have some kind of commitment to the idea that moral disagreements can be resolved by sufficient critical reflection which is why we bother to engage in moral debate" (137).

But commitment to moral debate is one thing, commitment to convergence is another. It would sufficiently explain the first commitment that we always hope to find a basis for agreement; that we expect that agreement will (or would) come seems like a hypothesis running far ahead of the data.

XI. Communication

A different reason for imputing a commitment to convergence has to do with moral communication. Suppose that we were to arrive not at "a single mature folk morality but rather different mature folk moralities for different groups in the community" (137). Then, Jackson thinks, "the adherents of the different mature folk moralities will mean something different by the moral vocabulary because [their] moral terms...will be located in significantly different networks" (137). One assumes that current speakers too, to the extent that they belong to precursors of Jackson's "different groups in the community," will have meant something different by their moral vocabulary.

But then, insofar as we see today's factions as in danger of evolving into "different groups in the community," we see them as in danger of not really communicating, on account of their similar-sounding words having different semantic values. Contraposing, we have got to see ourselves as converging if we want to see ourselves as communicating.

I don't know whether Jackson is attracted to this reasoning or not. But it looks to be in tension with a point he makes elsewhere in his Locke Lectures, in his discussion of color. He says that the folk might well take it for granted that there is "a kind, indeed a natural kind, distinctive of the exemplars of water and gold. [But] the folk are too sensible to have presupposed something as risky as that there is a distinctive kind in common to things we call 'red'" (108).□ The presupposition would be "risky" because it would put the legitimacy of our practice with "red" at the mercy of

developments that we're not in a position to prejudge. Presupposing moral convergence would be risky in the same sense. If the folk don't do it with "red," it stands to reason that they wouldn't do it with "right" either.

Then what does entitle moral disputants to regard themselves as communicating? Not enough attention has been paid to the cognitivist who answers like this: If something ought to be written into our folk theory, it's not that risky conditions ABC are met -- given which such and such a semantic mechanism sees to it that we attribute the same property by "right" -- but simply that you must mean something different by "right" is a diagnosis of the very last resort. The project of using evaluative language together in figuring out how to act is so overwhelmingly important that we do not allow others to opt out with this facile semantic excuse. Our primary commitment in other words is that "right" and similar action-guiding terms should stand for the same or similar properties in all our mouths.

The commitment comes at a price: a certain kind of anti-dogmatism. The readier we are to claim conceptual authority for our own evaluative views -- to say they follow from what "right" means in our mouths -- the harder it becomes to hold onto the idea of coreference as between disputing parties. But then, rather than its being a condition of moral communication that we expect to arrive at a single moral truth, the proper condition is that nothing will ever be regarded as the point of arrival: the point at which reference is finally fixed and moral theory acquires a conceptual imprimatur. This is what makes Jackson's second response to the open question so counterintuitive. He says that

the idea that what fits the bill that well might fail to be rightness is nothing more than a hangover from the platonist conception that the meaning of "right" is somehow a matter of its being mysteriously attached to the form of the right (152).

I would have thought it was part of the the bill that what satisfied it might still fail to be rightness. This is not because of platonism but the opposite: we refuse to attach "right" to any descriptive property so tightly that moral dissidents, even hypothetical ones, come out as simply misusing "right." Similarly, it seems part of what we have in mind by "reasonable" that the door is left open to the brilliant iconoclast who gets us to see that we have all along been acting contrary to reason. To dismiss such a person as meaning something else by "reasonable" strikes us as dogmatic. I doubt that as folk reason-theory matures it will take a kinder view of dogmatism than we do today. □

XII. Meanings as reason-laden

I have been saying that facts about what it's reasonable to say or do are one thing, facts about meaning another. As we've just been seeing, at least one modal rationalist denies this. He thinks that the meaning of 'reasonable' is enough to determine (with the help of lower level descriptive facts) its extension: what it is that's in fact reasonable. Now let's look at a second way in which the semantic and the epistemic get mixed up on Jackson's picture. Not only does the meaning of 'reasonable' close off all questions about its extension, its extension has an unexpected influence on the meanings of other words. The meanings of other words in my mouth are skewed by my sense of what's rational. Let me explain..

Suppose we have for each possible world v a complete physical specification P_v in physical terms; and suppose that (as Jackson maintains) my understanding of a sentence S is given by the truth values I assign to instances of the schema:

(S_{vw}) if P_v is true, then S would have been true had it been that P_w .

This applies in particular to statements S about mental states, including phenomenal mental states.

It seems clear, though, that one cannot determine the truth-value of phenomenal statements with respect to physically given worlds without asking what it would make sense to say about these worlds. What is the the best explanation of the fact that there is weeping and wailing and gnashing of teeth? That the weepers are in in pain, or that they are in such and such brain states but feel nothing? If you and I have sufficiently different ideas about this - - because of our larger theoretical and practical projects, how anxious we are to avoid multiplying entities, etc. -- we may be expected to assign different truth-values to S_{vw} . But then, just because of our methodological orientations, we wind up meaning different things by "there is pain"!

Another example: Suppose that you and I are confronted with a world in which events with the characteristic physical manifestations of pain occur on all the same occasions as events that we agree deserve to be called c-fiber-firings. You decide for ontological economy reasons that "pain" and "c-fiber firing" pick out one and the same type in this world. I am not too worried about ontological economy; I decide on the basis of my attachment to psychological autonomy and modal intuition that they pick out different though de facto correlated types. What does the Jacksonian physicalist say about this case? Again, that you and I mean different things by "pain."

That is not how ordinary speakers see it. Remember the great identity debates of the 1950s, when it was assumed that mental/physical correlations would soon be found and the question was what ontological conclusions to draw. The disputants didn't think of these debates as driven by differences about the meaning of "pain"; they thought they were arguing about the metaphysics of pain.

Of course, to repeat the obvious, everyone is entitled to use words however they like. But if Jackson is using "conceptual" in a special sense, to describe differences which others would classify as doctrinal, then that bears on the interpretation of his claim that physicalists are committed to the 'conceptual' entailment of the psychological by the physical. It's not clear that the claim comes to more than this: those who find it on balance reasonable to apply mentalistic description to a physically given world can portray those who disagree with them as "meaning" something different by the mentalistic description.

***** strong a priori vs. weak a priori

XIII. Physicalism

Now finally a word about the Jackson-Chalmers objection to physicalism. This objection turns in large part on the idea that there is no analytic or a priori entailment from our world's physical condition to its phenomenal condition. By thesis (A) in the first section, the failure of analytic entailment testifies to the existence of a world physically like ours but lacking in consciousness -- a so-called zombie world.

The modal rationalist chalks it up to my semantic competence with a word how I find it reasonable to apply that word in given circumstances. So it would be then with 'pain' as well. Now suppose that I find it reasonable to say of a world physically just like ours that there is lots of pain in it, and count the world into the primary intension of 'there is pain' as a result. Then it comes out analytic in my idiolect that if things are physically like so, there is pain. No failure of analytic entailment means no failure of analytic entailment to 'explain' by postulating a zombie world.

So to the extent that I was worried that I might have to give up my physicalism on account of the lack of an a priori link between physics and pain, I should stop worrying. The missing apriori link has been restored.

Am I saying that physics a priori necessitates pain, full stop? Not at all. The claim is rather that the modal rationalist faces a dilemma. If meanings are individuated so that people who disagree about what to think thereby mean different things, then for at least some of us, "if the world is physically like so, then there is pain" is true in virtue of (the a priori aspect) of meaning.

If on the other hand meanings are not to be individuated that finely, then the rationalist's claim that lower-level physics analytically necessitates everything but consciousness becomes hard to believe. It's directly at odds with a version of Quine's underdetermination thesis. And not a particularly controversial version. Positivism aside, it seems generally agreed that the methods we use to get from lower level data to high theory -- methods like inference to the best explanation, Ockham's razor, and so on -- are although perhaps rational not analytically valid. This is what the modal rationalist who picks the second horn of the dilemma must deny.

Appendix

It may seem that Chalmers has an easy answer: granted that there may be other ways of explaining a posteriority using worlds, he says he wants a world for every not apriori false hypothesis. (Analogy with Quine: worlds as not admitting of analytically inequivalent descriptions.) Each not a priori closeable gap between different ways of representing reality means a difference in how reality really could have been.

To answer Chalmers on his own terms, I need to show that it's not just that we can, we need to make room at some point for intentional differences that don't correspond to differences in the way that the world is represented as being, and so differences in which worlds make a given representation true. Here is the plan. Dave holds that

(A) whenever it fails to be a priori that all Ps are Qs, there is a world that construed as actual contains Ps which aren't Qs.

I want to explore a certain way of constructing counterexamples to this thesis, and then see if the method can be extended to encompass the case where P = physical stuff and Q = pain. The idea in a nutshell is that you will have counterexamples to (A) when P is an objective, third-personal predicate, and Q is a predicate that has to do with how things strike us perceptually: for short a seeing-as (or hearing-as, or smelling-as, or etc.) predicate.

A certain kind of line drawing will be seen by anyone who looks at it as a human face. I don't mean that everyone will judge it to resemble a face or to represent a face, just that we cannot easily stop ourselves from "seeing a face in it" and forming associated judgments, e.g., that the face looks cruel or alarmed or what have you. I will call line drawing like this "facial." Note carefully that faciality is not the dispositional property of looking like a face, it's the geometrical property whereby things do look that way.

So far so good. Now let the predicate P be: is such and such an arrangement of lines on a page. Let the predicate Q be: is facial. It is true, I take it, that all Ps are Qs. But it is not a priori true. To determine whether a P-type drawing is facial, I need to put one in front of me and see how it strikes me. According to principle (A), then, there needs to be a world that construed as actual contains drawings of the indicated kind that aren't facial.

You may say that such worlds exist. Imagine a world in which human neurophysiology is mixed up so that Ps are no longer the sort of thing we see faces in. Who knows, maybe our entire capacity for "seeing in" is taken away. A world like will be one such that if it is actual then not all Ps are facial; indeed none are.

So it seems that the modal rationalist has an answer. The problem is that this way of saving principle (A) makes trouble for another principle the modal rationalist accepts, viz. (A)'s converse

(B) whenever there is a world that construed as actual contains Ps which aren't Qs, it not a priori that all Ps are Qs.

To see why, consider an example of Chris Peacocke's:

...the difference between perceiving something as a square and perceiving it as a (regular) diamond is ... a difference in the way symmetries are perceived. When something is perceived as a diamond, the perceived symmetry is a symmetry about the bisectors of its angles [the diagonals]. When something is perceived as a square, the perceived symmetry is a symmetry about the bisectors of its sides. (TC, 76)

This seems to bear on our concepts of squareness and diamond-shapedness too. Both are concepts of geometrical properties, indeed the same geometrical property; it's just that in the one case the property is presented as that which makes us perceive an object as symmetry across one set of bisectors and in the other it's presented as the property whereby an object is seen as symmetrical about another pair of bisectors.

Now, on anybody's account, we know a priori that squares are diamond-shaped and vice versa. But now wait. Can't a world be found if actual make "squares are diamond-shaped" untrue? All we would need is a world in which due to a wiring mix-up we have lost the ability to see things as symmetrical about their diagonals.

Thus if (B) is to be at all plausible, we will have to interpret "there is world that construed as actual etc." as limited to worlds in which people perceive-as in the same sorts of ways they perceive-as here. Otherwise it will be all too easy to destroy bona fide a priori connections between seeing-as concepts by mucking inappropriately with "how we take things."

Another example. It's a priori that if option J is much more painful than K and they are otherwise the same, then the more painful option is worse. Our concept of "wrongness" ties it very closely with what strikes us as meriting disapprobation. But now imagine a world in which we for some reason consider the more painful option better.

This is a world which if actual has the a priori claim coming out false. Once again, then, in our applications of (A) and (B) we need to limit ourselves to worlds in which our semantically relevant dispositions are unchanged.

Actually one can say more. We should take our actual semantically relevant dispositions and bring them to bear on the world-imagined-as-actual. Chalmers in fact stresses this, to bring out the non-a-priority of "I exist and have thoughts." Thus in considering "all Ps are facical" we consider a world as actual and look at the Ps in that world and ask ourselves whether they are facical. The views of the inhabitants are neither here nor there since they may be nuts.

Consider in this light the diamond-square case. In considering "all diamonds are square" we consider a world as actual and look at the squares in that world and ask ourselves (not the inhabitants) whether they are diamond-shaped. It's what we mean by diamond-shaped, and so it's our semantic dispositions that matter. Since we are indeed disposed to see squares as diamond-shaped, the conditional comes out a priori.

All this seems just right. But now we have a problem for (A). We agreed above that it's not a priori that Ps are facical because we have to observe our own semantically relevant reactions to Ps. But our reaction is a given; we find them facical. It follows that there is no world which considered as actual contains Ps that aren't facical. Therefore claims employing grokking-as predicates can be non-apriori without a witnessing world.

Next step is to apply the moral to "all Ps feel pain." Suppose as an experiment that "pain" is a grokking-as predicate. It picks out its referent at least partly in terms of how we perceive that referent – what we invest it with. Then the following looks prima facie coherent. The referent that "pain" picks out is a physical property. The reason that it is not a priori that all Ps feel pain is not that there is a world which if actual contains painless Ps. The reason is that as with diamond-shaped, in evaluating P-worlds we have to consult our own reactions. Unfortunately, the relevant reactions are the reactions of people reacting to P "from the inside." (Something similar is true with "diamond-shaped" and "facical.") So you have to approach the world with an eye to how physical property P strikes you as an instantiator of P. Nothing for that but to do the experiment! (Just like facical.) You find it strikes you as hurtful.

Familiar point that you've got to hold something constant. What is it? Narrow meanings of words. But what if those narrow meanings are tied up with semantically relevant dispositions? Got to hold those constant too. In fact it's a mistake to talk about holding constant. What we do is bring our actual dispositions to bear on the counterfactual world considered as actual. At least those of the dispositions that are narrow meaning constituting.

Necker cube. Duck/rabbit. Duck/rabbit from different perspectives, or where you need to scrunch your eyes in different ways so you can't tell if figure has changed. Elaborate example so harder and harder to switch between the two perspectives.

Demonstrative reference given how little we know about what causal relations make for reference. "This is that." Maybe don't know enough about the mechanisms of reference in the two cases. It's possible for all you know that there are different referents but it may not be really possible. That depends on what the mechanisms of reference really are. Peacocke examples. "It's raining now" on different occasions when you've been told you'll be woken up multiple times.

Strategy of divide and conquer. I want to describe Two scenarios where "x is D iff x is E" is non-apriori without this meaning that there's a possible world with the D-feature and not the E-feature.

(I) Sometimes when "D iff E" is non-a-priori, that's to be explained by saying that there's a world which for all the speaker knows is D-ish without being E-ish. You have to leave that option open because to close it is to assume that speakers know essentially everything about what a referent has to be like to be the referent. And they don't. Example of this might be: it's not a priori that something this big is that big. (I'm referring to the sizes.) I doubt you need a world where the things are really of different lengths, it's enough that you can't tell they're the same length. That's enough to make it conceptually conceivable that the one is longer. Out one is: the non-transparency of reference-determination. (The lengths they looked to be were exactly the same, so there is no world where the one is represented but not the other.)

(II) Sometimes when "D iff E" is non-apriori, that's to be explained by saying there are centered worlds which are E-ish but not D-ish.

Example: I'm looking with both eyes through what might be the same tube or might be different tubes or who knows what. I rationally suspect that that (what I'm looking at with my left eye) isn't that (what I'm looking at with my right eye). I don't know anything to distinguish the relevant referential mechanisms but that they terminate in different eyes. What can we say about the world that "witnesses" the rationality of my suspicion? Only that it has me looking at distinct objects with my two eyes. What can we conclude about the properties of the referents? Only that the property of being looked at with this eye isn't the property of being looked at with that eye. There are worlds where things have the one "property" without the other.

Now try it with pain. I rationally suspect that that (pain) isn't that (c-fiber firings). I know next to nothing about the mechanism whereby 'pain' picks out pain and not something else: mainly just that pain viewed from the inside elicits in me a judgment of painfulness. What can we say about the world that witnesses the rationality of my suspicion that pain isn't c-fiber firings? Just that that world contains something with the property of eliciting the inside way a judgment of painfulness and nothing with the property of being a c-fiber firing. What can we conclude about the properties? The property of eliciting in me when viewed from the inside a judgment of painfulness is distinct from the property of eliciting in me when viewed from the outside a judgment of c-fiber-firingness. But no one ever doubted that those two were different properties. That's just the essential indexical point.

I don't expect you'll buy any of this, but suppose for the moment it's granted. A certain kind of dualism might still be arguable along the above lines. It might be arguable that the property of seeming-so-and-so-when-viewed-from-such-and-such-a-perspective is distinct from any "objective" property. This because it (almost) always makes sense to wonder how an objectively given thing will strike you when you confront it yourself. "I know it's so and so, but how does it look/feel/smell?" People might differ though about whether this marks a genuine metaphysical divide or is just one more manifestation of essential indexicality

Reply to Stephen White:

Suppose then that the belief (whichever one is supported in W by the relevant narrow content) does have to be true in W. Then the inference I'm not sure I buy is the one from "it's not irrational to believe that water is Q while H₂O isn't" to "there must be a world which justifies the belief" - ie., makes true whatever belief its narrow content induces in that world.

You say the following: "A world at which 'water' and 'H₂O' pick out different substances with the properties that Smith ascribes to them will be sufficient." A world like that probably does make Smith's narrow content true. But granted that such a world is sufficient for Smith's rationality, it might not be necessary. The door is thus left open to another sort of world, one not containing distinct substances with the properties Smith ascribes to them, doing the rationalizing job.

Now I realize this is a door you want to leave open, since you're interested in the possibility that Smith associates the same property with the two terms, but conceives that property first one way and then another. But I have something different in mind. You seem to be assuming that the mechanism by which terms S and T achieve reference to x and y involves

- (i) Smith associating properties A and B with the two terms, and
- (ii) x and y being the A- and B-things, respectively.

I think I agree if (i) is interpreted as

- (i') there being properties A and B such that Smith's term S (T) refers to z iff z is A (B).

The part I'm not clear about is why Smith has got to self-consciously associate the properties with the terms. Suppose for instance that S and T are perceptual demonstratives "that one" and "that other one." The only thing that marks x out as S's referent and y out as T's referent is a slight shift in Smith's attention as he switches from S to T. You can suppose if you like that x is slightly to the right of y in Smith's visual field, but that Smith for one reason or another attaches no representational importance to that. (He knows he's taking part in a visual-field-scrambling experiment.)

A slightly different caveat has to do with the nature of the properties A and B. It seems possible that A and B are not properties in the usual sense but perspectival and/or recognitional properties; some would speak of perspectival and recognitional concepts. So in the previous example, A is the property of being the thing I'm attending to now, or the leftmost thing in my visual field, or something like that. Or suppose that I'm looking at a Necker cube, that S and T are occurrences at successive moments of "the frontmost square now." Then A is the property of looking frontmost to me at such and such a time. One can think of lots of examples like this, e.g., "how big the room looks from here" vs. "how big it looked from there." So I am tempted to think that (i') needs to be something more like

- (i'') there being "properties" A and B such that etc.

where "properties" means properties in the loose sense in which perspectival/recognition properties are allowed, properties of the form: strikes me as so and so when I adopt such and such a perspective on it.

Now let me raise two worries about the semantic principle, one turning on the "association" issue raised two paragraphs and one turning on the "perspective" issue raised in the last paragraph.

(I) Smith might not know a whole lot about the mechanism by which the terms find their referents; or he might know a lot but not everything. He might not know everything either about the properties A and B that x and y had to have to be the referents of S and T. Suppose his knowledge of A and B suggests that A isn't B but leaves it open that A is B. Then it would be enough to save the belief that 'S isn't T' from irrationality if there was a world in which distinct things x' and y' had distinct properties A' and B' answering to the limited knowledge Smith possesses of A and B. It's compatible with all this that A is B.

The worry is that all that seems clearly needed to ground the posteriority of "S = T" is that a speaker is slightly underinformed about the properties x and y need to possess to qualify as the referents of terms S and T. No actual difference between the properties is required as long as there could be a difference as far as the thinker is in a position to tell. (It may be that allowing the speaker even this much ignorance of the mechanisms of reference prevents modes of presentation from playing the rationalizing role you have in mind for them.)

(II) Someone could say: oh, but speakers always know of at least one difference between the properties a thing has to have to be S's referent and the properties a thing has to have to be T's referent, viz. that to be the former you have to hook up with S in the right way, and to be the latter it's T you have to hook up with in that way. That is, x has the property of eliciting a judgment of S when viewed from such and such a perspective, and y has the property of eliciting a judgment of T when viewed or considered from such and such another perspective.

This seems fair enough, but I'm not sure how much it buys you in the way of a dualist conclusion.

Example: I'm looking with both eyes through what might be the same tube or might be different tubes or who knows what. I rationally suspect that that (what I'm looking at with my left eye) isn't that (what I'm looking at with my right eye). I don't know anything to distinguish the relevant referential mechanisms but that they terminate in different eyes. What can we say about the world that "witnesses" the rationality of my suspicion? Only that it has me looking at distinct

objects with my two eyes. What can we conclude about the properties of the referents? Only that the property of being looked at with this eye isn't the property of being looked at with that eye.

Now try it with pain. I rationally suspect that that (pain) isn't that (c-fiber firings). I know next to nothing about the mechanism whereby 'pain' picks out pain and not something else: mainly just that pain viewed from the inside elicits in me a judgment of painfulness. What can we say about the world that witnesses the rationality of my suspicion that pain isn't c-fiber firings? Just that that world contains something with the property of triggering 'pain'-tokens is distinct from that of being a c-fiber firing. What can we conclude about the properties of the referents? Only that the property of eliciting in me when viewed from the inside a judgment of painfulness is distinct from the property of eliciting in me when viewed from the outside a judgment of c-fiber-firingness. But no one ever doubted that those two were different properties. .

I don't expect you'll buy any of this, but suppose for the moment it's granted. A certain kind of dualism might still be arguable along the above lines. It might be arguable that the property of seeming-so-and-so-when-viewed-from-such-and-such-a-perspective is distinct from any "objective" property. This because it (almost) always makes sense to wonder how an objectively given thing will strike you when you confront it yourself. "I know it's so and so, but how does it look/feel/smell?" People might differ though about whether this marks a genuine metaphysical divide or is just one more manifestation of essential indexicality

Still looks to be a strong disanalogy between tables and consciousness vis a vis the lower level facts. One is tempted to say (and Dave does say) that the uncertainty with tables is due to unclarity about just what "table" means; it could be that "table" is vague and on way of resolving the vagueness tables definitely come whereas on another they don't. With consciousness however vagueness doesn't seem to play a role; it seems clearly conceivable that consciousness in the only sense that matters might be absent. I disagree that the table case is merely semantic but it remains I think that the consciousness case seems clearer. With consciousness it seems we can imagine it going either way without equivocation; with tables (mereological sums, etc.) it's not clear that can be done.

So I recognize a need to explain the disanalogy, and explain the (clearer) conceptual conceivability – non-apriority – of the physics/consciousness connection. The challenge is to explain how physics sans consciousness can be "rationally" conceivable without a world to witness the fact.

I've given the outlines of an answer elsewhere: the same world may be counted into the primary intension as fitted out with one center – as viewed from one perspective – and not counted in as fitted out with another. Now I want to expand on that.

Consider a minimal physical duplicate of this world: call it world Z. I say Z is the actual world, Dave disagrees; but he can't insist at this point that it's not, because I'm going to try to explain away on the basis of this assumption the evidence he cites against @ = Z. The evidence is roughly this. "If c-fibers are firing, there is pain" is a posteriori. So it is rationally deniable. So (here is the key step) there is a world in its primary intension, a world with c-fibers firing and no pain. This world would be Z.

As against this, I say that there is no world with c-fibers firing and no pain. What there is just world Z outfitted with different centers, from only some of which the pain appears.

A preliminary point: a centered world doesn't get counted into S's primary intension just on the surmise that S is true in it. If that were enough then the physicalist can say that he surmises there is pain. No, a world gets counted into S's primary intension just if someone operating from that center could have things stipulated to them as part of the presentation of that world that would decisively verify S.

Now we can agree that lower-level facts about a physically given world would not decisively verify "there isn't pain." So there is no way we can use any physically given world to witness "there is no pain." But the reason these worlds aren't suitable is of no metaphysical significance. It has to do with the problem of other minds.

world Z outfitted with some centers goes into the primary intension of "my c-fibers are firing" while Z outfitted with other centers goes into that of "I'm in pain." It is not the case that Z outfitted with any centers has "they're firing but I'm not in pain." This is because there is no single center from which you can mount both thoughts. "I am outerly presented with c-fiber firings" "I am innerly presented with pain."

One centered world based on Z has me here at the receiving end of phenomenal presentations. Call this ZI, for Z inner. Another centered world based on Z has a (sensationless?) someone looking at my brain states through a cerebroscope. Call this ZO, for Z outer.

If S is the sentence "I am in pain but there are no C-fibers firing," one can ask "is ZI in S's primary intension?" The answer is plausibly yes.