INTENTIONAL CHEMISTRY

Joseph LEVINE
North Carolina State University

1. In their book, *Holism: A Shopper's Guide*, Fodor and Lepore (F&L) argue for the minimal thesis that, despite its wide acceptance, no good arguments for the doctrine of Meaning Holism have been presented in the literature. This is a "minimal" thesis in the sense that they don't claim to show that there couldn't be any good arguments for holism, much less do they present positive arguments for competing doctrines. Yet it's hard to read the book and the article that followed (F&L 1992b) without the strong feeling that a positive agenda does exist here, and it's this positive agenda, and the uneasiness it seems to cause even in friendly circles, that I would now like to explore.

Since I'm mostly concerned with a debate "within the family", let's assume a Computational/Representational Theory of Mind. Mental processes are computational transformations of symbolic representations, and a mental state is a computational relation to a symbolic representation. The mind's representational system constitutes a "language of thought", with both a syntax and a semantics. The question at issue here, then, is the nature of the primitive semantic properties of "mentalese" terms. This is the computationalist form of the question of how to characterize the conceptual constituents of thought.

We can distinguish three positions with regard to at least one aspect of this question: atomism, molecularism, and holism. There are various ways of characterizing these positions, and they're not all equivalent, but I don't think the differences will matter for our discussion. Atomism is the view that none of a symbol's inferential connections to other symbols are constitutive of the conceptual content of that symbol. Molecularism is the view that some of them are, and holism is the view that all of them are. Of course the notion

of a "constitutive" relation is not altogether without problems, but the basic idea is this. If you think that I just don't count as believing that Jones is a bachelor unless I also believe that he's not married (and, consequently, have the concept MARRIAGE as well as the concept BACHELOR), then you think that the inferential connection between my mental symbol [bachelor] and my symbol [unmarried] is constitutive of the conceptual content of [bachelor].[1]

As I see it, molecularism is the common sense default position, abandoned only because it has been deemed indefensible. Holism and atomism have tended to be held as radical positions in response to the failure of molecularism. The figures cited in F&L's discussion all hold some version of holism, and a good part of the "positive agenda" I attributed to F&L above is to defend atomism.[2] In this paper I will not be concerned with holism, believing as I do that F&L have done it sufficient damage. Instead, I want to explore the question whether molecularism might not be capable of rehabilitation, especially within a naturalistic framework that takes a computationalist approach to the study of the mind. I will argue that the prospects for rehabilitation are not good, and therefore we ought to try to live with atomism. I'll also argue that this may not be as hard as it looks at first.

The main objection to molecularism is the one Devitt (forthcoming b) calls "the no-principled-basis" problem. Holists claim that all of a concept's inferential relations – "IR's" for short – are constitutive of the concept, or meaning-constituting (MC, for short). Hence the task of distinguishing those IR's that are MC from those that aren't isn't one the holist must undertake. However, since molecularists do make this distinction, on their view there must be some property that distinguishes the MC IR's from the rest.

According to the traditional pre-Quinean picture, the MC IR's are unproblematically identifiable with the analytic ones. On this picture, analytic truths are true by virtue of the meanings of their

1. I will use upper case to represent concepts, and enclosure in brackets to represent mentalese symbols. So, in these terms, the question at issue is whether any of the inferential connections that [bachelor] maintains with other symbols is constitutive of its expressing the concept BACHELOR.

2. Actually, to be fair to L of F&L, this positive agenda may belong only to F

terms, and they are known *a priori* because knowing them is an immediate consequence of competence with the relevant terms. If it's part of the meaning of "bachelor" that it applies only to unmarried men, then the sentence "all bachelors are unmarried men" is analytic and known (or knowable) *a priori*.

The problem for molecularism begins, then, with the Quinean attack on the analytic and the *a priori*. Quine argued that no sentence was immune to revision; that with enough imagination one could envision circumstances in which it would be rational not to believe it. But if one could give up the sentence without changing its meaning,[3] then the IR represented by this sentence couldn't be constitutive of its meaning. Hence no IR is MC.[4]

F&L state explicitly that they accept, or at least have no desire to challenge, the Quinean attack on the analytic and the *a priori*, so therefore they are not inclined to defend molecularism. This accords well with their (Fodor's, anyway) naturalism. But some philosophers, themselves of a naturalist bent, argue that F&L are too quick to grant the efficacy of the Quinean-based argument against molecularism. Since the two philosophers I have in mind in particular are Michael Devitt and Georges Rey, I'll call this naturalistic defense of molecularism the Maryland Defense.

Devitt and Rey adopt different strategies in their defense of molecularism. Devitt seems to accept the core of the Quinean attack on the *a priori*, but argues that this has no bearing on the prospects for a molecular theory of concepts. In particular, he thinks that a naturalistically acceptable notion of analyticity is possible without thereby bringing unacceptable *a priori* knowledge in its wake. By "unacceptable *a priori* knowledge" I mean knowledge based upon conceptual analysis and justifiable independently of experience.[5]

3. It's crucial that in the circumstances in question revision doesn't amount to a change of meaning. Otherwise, Quine's argument amounts to no more than the trivial claim that we could always decide to use our words differently than we now do.

4. See Quine (1953). From the characterization just given it looks as if Quine should be an atomist, but in fact, he's either to be read as a holist or a meaning eliminativist. See the discussion of Quine's position in F&L.

5. The qualification here is to allow for the possibility of *a priori* truths that are not analytic; for example, theorems of mathematics. Devitt probably doesn't

Rey goes further, arguing that the Quinean attack doesn't succeed even against this notion of the analytic-*a priori*. On the contrary, Rey claims that Quine's objections only work against a particularly superficial account of the analytic-*a priori*, one based on privileging introspectively-grounded meaning intuitions. Properly naturalized, in fact psychologized, there is an account of the analytic-*a priori* that is immune to Quine's objections. I'll take up each version of the Maryland Defense in turn.

2. Devitt's defense of molecularism proceeds in two stages. First, he argues that the Quinean attack on the analytic-*a priori* only affects theories that provide an epistemic criterion for distinguishing MC IR's from non-MC IR's. So long as a molecularist can provide a non-epistemic sorting of IR's, she will not run into trouble from Quine. Second, he argues that a suitable non-epistemic criterion is indeed available. In section 2.1 I will deal with the first stage, and in section 2.2 with the second.

2.1. Devitt notes that even F&L maintain that Quine's arguments tell only against theories that employ an epistemic criterion for determining whether or not an inferential relation is essential to a concept. He attributes their "unofficial" indictment of molecularism to their pessimism about the prospects for finding a suitable non-epistemic criterion. Indeed, it seems, from the passages he quotes (F&L, pages 56 and 57), that they agree with him that if such a criterion could be found, it wouldn't violate Quinean scruples against analytic-*a priori* knowledge. Whether or not this is what they intend, I will argue that such a position gives the molecularist too much. It's not just that there isn't a good non-epistemic criterion

like this sort of *a priori* either, but it seems to me that the status of mathematics is another matter.

In general, since all sides agree that any sort of meaning realism gives rise to synonymy, and therefore an analytic/synthetic distinction (see Boghossian, forthcoming), and since there might be *a priori* truths that do not result from conceptual analyses, as in mathematics, I will use the term "analytic-*a priori*" from now on to designate the category of statement with which I'm concerned. A statement has this status if its truth is guaranteed by a meaning analysis to which the subject has *a priori* access.

available, but that even if there were one, it would still have troubling epistemic consequences.

To begin with, let's try to clarify the distinction between employing an epistemic criterion and a non-epistemic criterion. One might hold that what determines which IR's are MC is their epistemic status. That is, a theorist would determine whether or not believing that bachelors are unmarried is constitutive of having the concept BACHELOR by first determining whether the belief is analytic-*a priori*. This would clearly count as employing an epistemic criterion to sort the MC IR's from the non-MC IR's.

On the other hand, one might hold that some feature of an IR other than its epistemic status is responsible for its being MC, and this would be to use a non-epistemic criterion. Suppose, for example, that the feature in question is some aspect of its causal role. On this view, we don't first look at an IR's epistemic status and then use this to determine whether or not it's MC. Thus, at least in advance, we don't need to determine that some beliefs are analytic-*a priori* in order to discover which of a concept's IR's are MC.

As just characterized, the difference between using an epistemic criterion and a non-epistemic one is in fact a real difference. But it's not that the non-epistemic criterion allows you to avoid commitment to the analytic-*a priori*. Rather, the difference between the epistemic criterion and the non-epistemic criterion has to do with the direction of dependency between being MC and being analytic-*a priori* – for the former, being MC is dependent on being analytic-*a priori*, whereas for the latter, being analytic-*a priori* is dependent on being MC. But either way, you're stuck with the analytic-*a priori*.

The argument for this last claim is fairly simple. I take the following two properties to be definitive of analytic-*a priori* belief: (1) that it cannot be empirically disconfirmed; and (2) that it cannot be given up without change of meaning (or concept). These two features are of course connected. If you can't revise a belief without change of meaning, then you can't empirically disconfirm it. For suppose some empirical evidence, perhaps having to do with overall theoretical cogency, convinced you to give up the relevant belief. Since, by hypothesis, to give it up entails changing its meaning, it

isn't *that very belief* that you've given up. Hence, it can't be empirically disconfirmed.

Note that it's just these two features that Quine's attack on the analytic-*a priori* is aimed at. He argues that with enough adjustment in our belief system as a whole, any single belief could be given up. He's not merely pointing out that we could decide to stop saying "all bachelors are unmarried", for even the defender of the analytic-*a priori* admits that. Rather, he's arguing that, in response to sufficient empirical pressure, we could stop saying it without changing its meaning.

Now suppose, using some allegedly non-epistemic criterion, we discover that believing all bachelors are unmarried is constitutive of having the concept BACHELOR. Then it follows that one can't cease to believe[6] that all bachelors are unmarried without losing the concept; or, as we might say, changing the meaning of the relevant term. But then that makes the belief that all bachelors are unmarried analytic-*a priori*, which is what we needed to show.

Devitt tries to meet this sort of objection in two ways. First, he argues that though the molecularist is committed to the claim that competence with a concept entails having a certain belief[7], she is not committed to the claim that "... competence alone can supply the justification that turns the belief into *knowledge*".[8] Second,

6. Of course, as Georges Rey pointed out in conversation, there are subtleties here about what constitutes ceasing to believe something. After all, one can stop saying "all bachelors are unmarried", or even explicitly deny it, without actually ceasing to believe it; perhaps one is confused, or even maintaining contradictory beliefs. My point still stands, however. Take whatever epistemic state it is that is supposed to play this constitutive role, losing that state entails losing the concept.

7. Devitt distinguishes between what he calls the "inferential version" and the "belief version" of molecularism. On the former, it's the tendency to infer from [x is a bachelor] to [x is unmarried] that's MC, whereas on the latter it's the explicit holding of the belief [all bachelors are unmarried] that does it. He claims that the sort of problem I'm raising here only arises, even *prima facie*, for the belief version. But since one can easily define a notion of *a priori* inference that would do all the work of *a priori* belief, the distinction doesn't seem relevant to me, and I will ignore it in what follows.

8. This quote, and the following 2, are from sections III.4 and III.5 of Devitt (forthcoming b).

though he admits that a certain sort of unrevisability follows from the molecularist position, he claims that it has no epistemic significance; it is "*harmlessly metaphysical*".

With regard to the first point, Devitt takes the problem with the analytic-*a priori* to be its alleged ability to yield knowledge independent of experience. Since it's an empirical question, on his view, which of a concept's inferential liaisons are essential, from the fact that, say, UNMARRIED bears an MC relation to BACHELOR, it doesn't follow that one can know that bachelors are unmarried independently of experience. After all, you'd have to know the empirical semantic theory first. As he puts it, the knowledge acquirable by mere competence with the relevant concept "... does not supply any basis for foot-stamping dogmatism about it ... [it] supplies no privileged knowledge."

Well, it depends on what you mean by "privileged knowledge" here. If "privileged" entails immediate, conscious, introspective access, which is what I assume Devitt means by "foot-stamping dogmatism", then I agree that his view does not entail that we have "privileged knowledge". If this were all the Quinean naturalist were concerned with, then perhaps that would be the end of the matter.

However, I don't think this is the end of the matter, since it seems to me that there is another aspect of the attack on "privileged knowledge" that goes deeper than mere "foot-stamping". In this deeper sense, a belief is "privileged", or analytic-*a priori*, to the extent that it cannot be disconfirmed by empirical data (which includes broadly theoretical considerations as well). If indeed UNMARRIED is constitutive of BACHELOR, then the belief that all bachelors are unmarried cannot be empirically disconfirmed, as argued above. It doesn't matter that we may not know *this fact* – i.e. that [all bachelors are unmarried] is analytic-*a priori* – *a priori*.[9]

9. Notice that the claim that UNMARRIED is constitutive of BACHELOR does not entail the claim that the two second-order concepts CONCEPT OF UNMARRIED and CONCEPT OF BACHELOR stand in the same relation. Put another way, as Devitt himself points out by way of Morton White, the claim that "all bachelors are unmarried" is analytic does not entail that "'all bachelors are unmarried' is analytic" is analytic. I suspect that Devitt, though recognizing the level distinction here, is nevertheless smearing it somewhat. That is, he supports the claim of non-privileged epistemic character of the object-level belief by ad-

It is still the case that the relation between being a bachelor and being unmarried is knowable independently of experience, and that is something the Quinean naturalist doesn't want to accept.

Similar considerations apply to Devitt's second argument, that the brand of unrevisability you get from using a non-epistemic criterion to sort the MC IR's from the others, has only metaphysical consequences, not epistemological ones. He supports his case by drawing an analogy between meanings and other theoretical entities. Being a planet, or a capitalist, is a genuine property of objects, and it bears the constitution relation to other properties as a matter of fact, not decision. Just as there need not be an analytic-*a priori* definition (i.e. specification of the "planet-constituting astronomical properties") of "planet" available in order to justify the claim that not every astronomical property of a planet is essential to its being a planet, there need not be an analytic-*a priori* definition (i.e. specification of the MC IR's) of BACHELOR in order to justify the claim that not every IR is essential to its identity.

We can put the matter this way. Someone might think it's a matter for meaning-analysis to discover what is essential to being a planet, and then Quine's objections would intrude. If we take the discovery of the essence of planethood to be an empirical matter, however, then Quine cannot object. Devitt seems to think that the same move is available to the molecularist; eschew meaning-analysis in favor of empirical investigation.

The problem with the analogy is that the difference in subject matter here – planets vs. meanings – is crucial to the question of whether or not epistemological consequences follow from the discovery of essential properties. Unlike the case of planets, in the case of meaning, by its very nature, you can't avoid meaning-analysis, for that is what an empirical investigation of the sort Devitt recommends will yield. Of course the procedure for determining what the essential properties are (i.e. which IR's are MC) is empirical. But in the end, once the essential properties are determined, what you have is a meaning-analysis for the concept in question; and it is precisely the existence of such analyses that generates the sort of

verting to the admittedly empirical character of the meta-level belief. It seems to me that the latter does not provide support for the former.

privileged epistemic status to which the Quinean naturalist objects.

Let me elaborate this point with an example. Suppose that we're investigating the concept CAT (just to vary the example for a change). According to the traditional, pre-Quinean method, we would consult our linguistic intuitions and thereby determine which IR's were essential. So, we might claim that only animate objects could count as cats. Clearly, this method would result in claims of analytic-*a priori* status for [all cats are animate]. Just as clearly, Quine's objections to this status would apply.

Now along comes Devitt and argues that we don't consult linguistic intuitions but rather conduct empirical psychological investigation. Okay, let's say we discover thereby that the inference from CAT to ANIMATE is essential to CAT.[10] This discovery is apparently on a par with the discovery that revolving around a hot ball of gas is essential to being a planet. In fact, though, there's a crucial difference. While the latter discovery does not result in analytic-*a priori* status for the sentence [All planets revolve around a hot ball of gas], the former, though arrived at empirically, does result in analytic-*a priori* status for the sentence [All cats are animate]. For if one failed to accept this sentence, given the MC status of the inference from CAT to ANIMATE, one wouldn't mean CAT by [cat] and therefore the sentence can't be given up without change of meaning; precisely what it is to be analytic-*a priori*.

To sum up, I grant that one can distinguish between a position that looks to epistemic status as the basis on which to determine whether or not an IR is MC and one that doesn't. But even the latter position will entail significant epistemic consequences, since, an IR, by its very nature, is an *epistemic* property. Hence, molecularists are committed to the analytic-*a priori*.

Let me conclude this discussion on an exegetical note. I've been interpreting Devitt all along as someone who wants to avoid commitment to the analytic-*a priori*. Perhaps this is not correct. It might be that in fact it's purely the issue of introspective access that

10. One might wonder how empirical investigation could yield the result that such a connection is *essential* to CAT, as opposed to merely holding for all possessors of the concept. This is in fact a difficult question. A good part of my argument against the Maryland Defense is that there isn't a good answer available.

concerns him. If so, then, for my purposes, his position collapses into Rey's. My principal purpose here has been to establish that molecularism is indeed saddled with an epistemic doctrine which the naturalist tradition emanating from Quine has been unwilling to accept. Now, of course it's open to the molecularist to challenge forthrightly this unwillingness, and I will explore that sort of challenge in section 3. But first let's deal with the second stage of Devitt's argument; that is, his claim that a suitable non-epistemic criterion is in fact available.

2.2. Before dealing directly with Devitt's proposed criterion, some remarks are in order regarding his discussion of the need for such a criterion. Devitt argues that there are two sorts of questions one might be asking when demanding a principled basis for distinguishing the MC IR's from the others, only one of which really requires a substantive answer. He calls them the "descriptive" question and the "normative" question.

The descriptive question is a question about the constitutive properties of the "... *meanings we do as a matter of fact ascribe* ..." in "that" clauses (Devitt, forthcoming a, section II, emphasis in original). That is, we have a practice of ascribing beliefs to others, and, as a matter of fact, in doing so we take some, but not other, beliefs (or inferential relations) to be essential to the identity of the ascribed belief. From this practice we can abstract a meaning that is ascribed, and it's just a matter of fact which of its properties are constitutive and which aren't. No special criterion is needed.

On the other hand, the normative question is a question about the constitutive properties of the "... meanings *we should ascribe for semantic purposes* ..." (*ibid.*, emphasis in original). There are many possible meanings that could be attached to any ascription via a "that" clause. I say that Jones believes that cats are lazy. The term "cat" can be taken to express any of a large set of possible meanings, where each meaning is identified with a set of IR's (or associated descriptions/beliefs). The descriptive issue is resolved by discovering which of these possible meanings I am in fact ascribing to Jones. But the normative issue is not so easily resolved, since we want to know, given certain semantic purposes, which of the possible meanings I ought to ascribe. Devitt admits that to resolve *this* question a

principled criterion is called for.

While I have qualms about even the descriptive issue,[11] it seems to me that my concern in this paper falls squarely on the side of the normative issue. Devitt's talk of "semantic purposes" may not correspond precisely to the goal of producing a principle of individuation for intentionally characterized psychological states, yet it does seem clear that when we are talking about individuating entities to which we appeal in the context of theoretical explanations, we are in the normative realm. At any rate, in so far as Devitt admits there is some purpose for which a principled basis for sorting MC IR's from non-MC IR's is called for, I want to investigate whether or not the one he proposes can do the job.

Basically, Devitt claims that an IR is MC if it determines reference (*ibid*). The rationale for this criterion derives from what he calls the "Fregean Assumption", the principle that meaning is a mode of presentation, a mediator between the mind and the world. On this conception of meaning, it makes sense that those properties which aid significantly in the fulfillment of this mediator role should count as essential to meaning.

However, I don't think the criterion works. The problem has to do with what is meant by "determining reference". I see several possible meanings, three of which I would group under the rubric of what I will call "satisfactual determination". I also think there's another sense I will call "causal determination". My argument is that on only one of these readings is it plausible to pin being MC

11. My concern is this. If I understand him correctly, Devitt's characterization of the descriptive issue seems to take it for granted that there is a systematic character to our practice of belief-ascription. However, suppose that no two people ever ascribe the same meaning (i.e. set of IR's) when using the term "cat" in a "that" clause, or that even a single person doesn't ascribe the same meaning from one occasion to another. If this were the case, then I think we would cease to talk of meaning-ascriptions altogether, but rather of pragmatic considerations that govern which terms are employed in an ascription, or which substitutions are allowed on different occasions.

Of course the atomist can allow for this possibility, since all she needs to secure sameness of belief is sameness of reference. That is, whatever stability there is to our ascriptive practice may be accounted for by reference, and whatever instability there is may be chalked up to pragmatics.

on determining reference, but, unfortunately, on that reading the criterion turns out to be circular, and therefore useless.

Let me exemplify these distinctions with a concrete example. Suppose the question of meaning constitution concerns the term [cat], as it behaves in Sally's language of thought. In Sally, [cat] possesses many IR's, among which is its relation to [animate]. By saying the IR to [animate] satisfactually determines the reference of [cat], I mean that for any x, [cat] refers to x only if [animate] refers to x. But having said this, there are still various ways of reading this condition.

First, one might intend merely a material conditional – as a matter of fact only animate things are cats. Clearly this is not strong enough, since this would make all of Sally's true beliefs about cats constitutive of her concept. Second, one might intend the conditional to hold in all possible worlds, so one is saying that necessarily for all x [cat] refers to x only if [animate] does too. But this is clearly too strong, since there are bound to be possible worlds where [cat] refers to things outside the extension of [animate] (e.g. mountains).[12]

Clearly we want a necessity claim here, and perhaps this third one will do the job. Take the extensions of [cat] and [animate] in this world, call them "F" and "G". Then, necessarily all F's are G's. Of course we could just as easily have said this: necessarily all cats are animate. In other words, this third way of expressing the satisfactual determination condition amounts to saying that being animate is necessary for being a cat.

Of course this won't work either. The question of which IR's are essential properties of the meanings of symbols is not reducible to a question about the essential properties of their referents. Such a reduction would serve no psychological semantic purpose, besides being just plain counterintuitive. After all, if we followed this rule, then descriptions of genetic structure would be constitutive of conceptual content. Rather, we must separate the question of the role of a symbol in determining the referent of [cat] from the question of which of a cat's properties are essential to its being a cat. Many of the properties to which we would appeal to answer the

12. That is, take a possible world where [cat] refers to inanimate things and [animate] still refers to animate things.

second question (e.g. its genetic structure) will not be relevant to answering the first.

I see only one possibility left. To say that [animate] (satisfactually) determines the reference of [cat] is to say that there is a linguistic rule to the effect that [cat] refers to x only if [animate] refers to x, or by definition [cat] refers ..., or, finally, that it's an analytic entailment that if [cat] refers to x then [animate] refers to x. Of course these are all ways of saying the same thing; namely, that [animate] constitutes part of the meaning of [cat], and therefore the proposed criterion is blatantly circular. We are back where we started, in search of a criterion for distinguishing the MC IR's from the others.[13]

Perhaps instead of satisfactual determination, we should understand Devitt's proposal in terms of causal determination.[14] What I mean by "causal determination" is this. Let R be whatever the reference relation turns out to be. That is, [cat] refers to whatever it bears R to. Perhaps R is itself a causal/informational relation. To make things simpler let's suppose that's the case. An IR helps to causally determine reference, then, if it plays a significant causal role in maintaining the relation R between the concept (i.e. the symbol that expresses the concept) and its referent. Now there is no threat of a circle since we have the relation R to pin down the reference independently – i.e. conceptually independently – of the IR itself.

For example, suppose [cat] refers to the property of being a cat in virtue of their standing in the appropriate causal/informational relation to each other. It very well might be that our [cat] symbol would fail to stand in that relation to cathood were it not for the associated belief that cats are animate. In that case, the IR from CAT to ANIMATE would count as MC.

But if this is what Devitt has in mind, it doesn't seem very

13. It's important to emphasize that we are talking here about the normative question, which Devitt explicitly acknowledges requires an answer. Of course if we were discussing the descriptive question, he could say that no criterion is needed, so we aren't in search of anything. (But see note 11.)

14. In a personal communication Devitt informs me that he doesn't have causal determination in mind. Nevertheless, I think it's worth exploring.

plausible. First of all, it isn't at all clear what the motivation is for making the ontological leap from the discovery that an IR plays an essential causal role in establishing/maintaining the referential link to considering its playing that role to be constitutive of the term's meaning. What theoretical gain is there? After all, the real theoretical work is being accomplished by the referential relation R itself. It's the fact that the referent stands in the relation R that seems to *make it* – ontologically speaking – the referent, and the IR's role only comes in by way of helping to maintain that relation.

Put it this way, if the only reason we have for considering the IR to be MC is that it helps to maintain/establish the relation R between the term and its referent, why isn't the fact that R is established enough to constitute the referent as the term's meaning? If there's more to the justification of a certain IR's being MC than merely its playing this helpful role in maintaining/establishing the relation R between term and referent, then we must mean something else by "determine reference", like "satisfactually determine", and then we're back to the criticisms that attach to that theory.

Furthermore, and this is related to the point above, suppose we find that there are many ways for the crucial relation R to be established/maintained. In one case, IR1 plays an essential causal role, but in another case IR2 does. Sometimes IR1 can be replaced by IR2. Is there any basis for saying that the meaning changes when IR1 is replaced by IR2, or that for two terms that bear R to the same referent, but one does it by way of IR1 and the other by way of IR2 that they don't mean the same? I don't see what it is, at least not if we're concerned merely with the determination of reference. Of course there are other concerns that motivate IR theories, having to do with explanatory roles in psychology, etc., but I don't see how these have anything to do with a criterion of meaning identity that makes an IR MC if it plays a causal role in maintaining relation R.

Finally, there are threats of holism-type problems here as well. For all we know a whole lot of beliefs are causally necessary for the maintenance of relation R between a single term and its referent, yet it would seem quite odd to make them all MC. In fact, perhaps some false beliefs are part of the causal package. Do we want to make false beliefs MC too? Again, the point is that once we have a causal relation to constitute the referential relation, we can ignore

all those psychological states that are causally necessary for its maintenance when answering the ontological question of what constitutes meaning.

I think the main point that emerges from our discussion of Devitt's defense of molecularism is the inevitability of facing the question of the analytic-*a priori* if one is to mount a convincing defense. We saw earlier that no matter what criterion one uses to sort the MC IR's from the others – even one that does not explicitly appeal to the *a priori*, like the one Devitt proposed – the result is to make some sentences analytic-*a priori*. Furthermore, without appeal to the *a priori* in the first place, it is hard to see what motivation there is for making any candidate condition on IR's an MC condition.[15] So it's time to stop taking the Quinean attack on the analytic-*a priori* for granted, and see whether or not a suitably revised notion of the analytic-*a priori* might be acceptable and exploitable for the molecularist's purposes.

3. Rey (forthcoming) quite explicitly attempts to resurrect the analytic-*a priori* within a naturalistic – indeed a computationalist – framework. His principal defense of the analytic-*a priori* against Quinean attacks is based on a distinction between the claim that mere competence with a concept automatically yields endorsement come what may for certain sentences and the more modest claim that some sentences are the output of a "pure" cognitive activity under ideal conditions. The traditional pre-Quinean doctrine of the analytic-*a priori* was committed to the first claim, and it was to the point for Quine to note that almost any sentence could be accepted/rejected so long as sufficient adjustments were made elsewhere in one's theory of the world. According to Rey, all Quine's point comes to is the unsurprising recognition of human fallibility. We can think we know something *a priori*, and yet be wrong.[16]

15. This might be the reason F&L are so skeptical of finding a non-epistemic criterion in the first place. The point is that it looks like only an IR's playing a certain epistemic role could ground a claim to being MC.

16. As Rey (forthcoming) puts it: "Now, it can seem to follow from [Quine's claim about revisability of any sentence] alone that nothing could be analytic, since any two concepts could potentially be dissociated in this way ... . But this would be a mistake ... . We need to distinguish a dissociation of concepts in occurrent *belief* from one in *fact*." (Emphasis in original.)

While eschewing the position that mere competence with a concept yields immediate access to analytic-*a priori* truth, Rey claims that there is still a place for this category of truth in a theory of concepts. Basically, he argues that our empirical theory of the mind might reveal that certain cognitive processes, when abstracted from interfering factors such as stress, time and space limitations, and even other beliefs[17], result in *a priori* rules[18] that govern the application of terms; i.e. rules which, under ideal conditions, would be followed no matter what the sensory evidence.

Once we have a working notion of the *a priori*, it is possible to introduce a notion of the analytic in terms of it. Analytic truths will be the *a priori* ones, or at least that subset of the *a priori* ones that involve conceptual analyses. So, for instance, if our cognitive architecture is so organized so that, under ideal conditions, we would only apply the term [bachelor] to unmarried men, it is plausibly attributable to the fact that UNMARRIED is part of the analysis of BACHELOR. With conceptual analyses in play, we have a molecular theory of concepts.

According to Rey, the Quinean attack on the pre-Quinean tradition really only touches the "transparency thesis"; the view that analytic-*a priori* truths are determinable by introspection alone.[19] I will argue, however, that the Quinean attack goes deeper than that,

17. The idea here being that someone may hold a fairly bizarre theory which causes them to contradict some analytic-*a priori* truth. This aspect of the idealization was emphasized by Rey in conversation. Of course, in the end it seems to me that one has to be able to eliminate reference to other beliefs in the characterization of the ideal conditions. For under ideal conditions one presumably wouldn't believe such bizarre theories. If this weren't so, then what makes the conditions ideal? Anyway, I will not pursue the question of how to characterize the idealization in this paper, though this is clearly a central question for a view like Rey's.

18. I speak here of "rules" rather than "beliefs" because Rey, in conversation, has insisted that he is not committed to a representation of the rule constituting a full-fledged belief. It's enough if such a representation controls our application of the term in question, perhaps by occupying a special "file" for such representations. I will have more to say about "files" below.

19. "... they might assume that meaning is *transparent*, i.e. that the conditions on semantic competence require an agent to know a concept's analysis." Rey (forthcoming).

posing a serious challenge to any account of the analytic-*a priori* – even an idealized, psychologized one of the sort Rey proposes.[20]

3.1. I will begin by laying out some metaphysical and epistemological assumptions that will guide the discussion. I take properties to be objective features of the world that stand in accidental, nomic, and metaphysically necessary relations. That all the coins in my pocket at the moment are pennies is an example of the first sort; that water boils at 212° F. is an example of the second sort; and that H_2O contains hydrogen is an example of the last sort. Furthermore, these metaphysical relations are not automatically reflected in epistemological ones. Thus, if water is indeed identical to H_2O, then it is metaphysically necessary that water contains hydrogen. Yet, that water contains hydrogen – at least under that description – is not knowable *a priori*; we wouldn't want our theory of the *a priori* to entail that it was. If molecularism is not to slide into holism, it is precisely statements like [water contains hydrogen] which must be preserved from analytic-*a priori* status.

The challenge for the molecularist, then, is to distinguish cases of necessary relations among properties that reflect necessary relations among concepts from those that don't, and this can't be done from the side of the properties themselves. From the objective, or metaphysical point of view, there is no real difference between the sense in which water necessarily contains hydrogen and the sense in which being a bachelor necessitates being unmarried. If we are to find a difference in the mode of relation involved, it is going to have to be an epistemic difference, and this sort of difference must be grounded in the concepts involved – the modes in which these properties are presented to the mind.

As an example of two sentences that differ in epistemic status though not in metaphysical status, take the following sentences:

20. An important caveat is in order here. The published work of Rey's to which I refer in this paper is a very short piece, and therefore the view is only sketched in the barest outlines. There is a longer work, that is as yet unpublished, in which more details are filled in; but it is still a work in progress. My intention here is not so much to criticize Rey's position specifically, but to use it to exemplify my view about the general prospects for naturalizing the analytic-*a priori*.

[water contains hydrogen] and [H_2O contains hydrogen]. They both represent the very same metaphysically necessary situation, yet the latter is arguably analytic-*a priori* while the former clearly isn't. The difference between the two cases seems to be that the latter is true in virtue of its logico-syntactic form, whereas the former isn't. So, one clear criterion for the analytic-*a priori* that suggests itself is this: a statement is analytic-*a priori* if it's true in virtue of its logico-syntactic form.

Of course Quine claimed that even this was not a principled criterion, but none of his arguments in "Two Dogmas" directly addressed this criterion. Where he did address logical truth directly, his argument against according it a special epistemic status depended on the lack of an epistemically grounded canonical notation. If the choice of canonical notation is open, then of course what would count as a truth of logic would be open as well.

However, here I do think the computationalist, who is a realist about mental representation, has an answer to Quine. If indeed there is a medium of mental representation, then it makes sense to suppose it has a canonical format/notation at the level at which it is semantically evaluable, and that this is a matter of empirical, psychological fact. If so, there would seem to be a principle for determining which mental representations are true by virtue of logical form and which aren't.

If we can indeed make this distinction in a principled manner, then we certainly have one basis for arguing that [all bachelors are unmarried] has a different status than [water = H_2O] – the former involves a real decomposition into a representation that has a logically valid form, whereas the latter doesn't. Nothing in F&L's arguments oppose this. Quite the contrary, they claim that the only principled sort of (epistemic) analyticity there is is the one that derives from the compositional structure of the language; e.g. [brown cows are brown]. There is of course another view of Fodor's according to which there aren't any lexical decompositions of this sort (or very few), but this is an empirical matter that is orthogonal to the question at issue.[21] If you think [bachelors are unmarried] has

21. For extensive argument against lexical decomposition, see Fodor (1981) and Fodor, Garrett, Walker, and Parkes (1980).

a different status from [water = H_2O], one way to account for it is by way of lexical decomposition.

It's important to note, however, that, so construed, the analytic-*a priori* status of [bachelors are unmarried] is relativized to the subject; i.e. it's a feature of the subject's representation, not of the proposition represented. For instance, on this way of looking at it, there could be two subjects, both of whom are thinking the thought that bachelors are unmarried – i.e. a thought with that (wide) content – and yet for one it is *a priori* because true in virtue of logical form and for the other it isn't. This could happen because the very same property, being a bachelor, can be represented primitively by one subject and compositionally by another.

The same is of course true of [brown cow] as well. [Brown cows are brown] is logically/*a priori* true. But now let's suppose some subject has a primitive term [b-cow], the referent of which is brown cows. For this subject, [b-cows are brown] is not a logical truth. Rather, it has the same status as [water contains hydrogen] – necessarily true, but not *a priori*. (This just reiterates the point above concerning the distinction between the objective/metaphysical situation and the epistemic situation. It is representations that are compound or primitive, not properties.[22] Therefore it is only with respect to the former that the issue of truth in virtue of form applies.)

The sort of analytic-*a priori* we have just acknowledged – the explicitly compositional variety – is not really germane to the debate between atomists and molecularists. After all, atomism is a doctrine that is intended to apply only to symbolic primitives – atoms. Also, the problem for molecularism we've been investigating – how to characterize a principled basis for distinguishing the MC IR's from the rest – is clearly not a problem for compound representations, since the compositional rules of the language provide such a principle straightforwardly. Therefore, the debate at issue revolves

22. Actually, I don't know what to say about properties; perhaps there are compound properties as well. In fact, it may be with properties as well as representations that their necessary relations are to be exhaustively accounted for in terms of their compositional relations. Either way, the point at issue here is that the compositional relations among representations are not direct reflections of such relations among properties.

around the case of representational primitives. Supposing that one's concept of bachelorhood is represented primitively – as [bachelor] – is there a way to distinguish the epistemic status of the judgment that all bachelors are unmarried from that of the judgment that water contains hydrogen, given that both are metaphysically necessary truths?

By assumption, we've ruled out an externally-based criterion in terms of necessary relations among the properties themselves. Therefore, I see two options left. Either find a purely internal, psychological difference between the analytic-*a priori* judgments and the rest, or ground the distinction on a combination of internal and external factors. I will explore each in turn.

3.2. One suggestion mentioned by Rey is that associated with the mentalese terms that express concepts are "files" containing significant categories of information concerning the concepts involved. One piece of information might involve the canonical method of verifying the presence of instances of the concept. But there might also be information that is treated by the subject as essential to the concept; perhaps in some functionally specifiable way this file entry serves as a rule governing the subject's use of the term. On this story, to analyze a concept is to avail oneself of information stored in the relevant file, and this information would thereby achieve analytic-*a priori* status. Two individuals whose files differed in this respect would not be said to have the same concept, even if the associated symbols represented the same property. This account, then, is an example of what I'm calling a purely internal account.

It's crucial to this view that the contents of the relevant files are ultimately an empirical matter, or at least not immediately accessible merely in virtue of competence. Rather, the content of such a file is to be revealed both by philosophical reflection and by psychological investigation into the architecture of the mind. This view thus allows for the fallibility of linguistic intuitions of the sort Quine makes so much of, without giving up entirely on the notion that concepts do in the end have analyses.

It seems to me, however, that the file story faces a dilemma. Either it reduces to the compositional story discussed above, or the privileged epistemic status accorded to the file entries can't be justified.

The first horn of the dilemma is straightforward. One way to ground the claim of *a priori* status for the information contained in the file is to treat the symbol associated with the file as an abbreviation for it, or as an index for (or pointer to) it. What this amounts to computationally, is that when one reasons with the symbol one is actually reasoning with the contents of the file. But then we are really dealing with analyticity based on compositionality, and we haven't developed a molecular theory of primitives, which is what we were after.

So let's suppose that we don't go this route. It's not that the symbol [bachelor] abbreviates, or indexes the file containing [unmarried], but rather that there just is a file for [bachelor] containing the information that anything to which [bachelor] applies is also something to which [unmarried] applies. Here is where Quine comes along and asks: what endows this information, true (and maybe even necessary[23]) as it is, with the status of the *a priori*? Is it because it is found in a box entitled "essential properties file"? Quine's challenge – put in the terms of computational psychology – is to provide a functional/computational difference between such *a priori* belief and belief that's merely held very strongly and for good reason.

Obviously, one can't make the distinction just by labelling the file "concept analysis", since what we want to know is how anything would change if we just relabelled it "centrally held beliefs". But now we come to the idealization mentioned earlier. Suppose it is a fact about us, that no matter what evidence we are presented with, we never would (under suitably idealized epistemic conditions) give up the entry [unmarried] in the [bachelor] file. Thus, in the relevant sense, the belief that bachelors are unmarried would be unrevisable, and therefore analytic-*a priori*.

Here, then, is the core difference between the two beliefs – both involving necessarily true propositions – between which we need to make an epistemic distinction: (1) water contains hydrogen, and (2) bachelors are unmarried. Presumably, the story is that there is evidence, or possible evidence, which, even under suitably idealized

23. Of course Quine wouldn't himself admit the category of necessary truth. However, I do, so I'm willing to give this to the defender of the analytic-*a priori*. My point is that it still doesn't help.

epistemic conditions, would cause us to deny that water contains hydrogen, whereas there is no such evidence that, under the same idealized conditions, would cause us to deny that bachelors are unmarried. But now we need to know what sort of fact this is about us that we wouldn't ever deny that bachelors are unmarried.

One possibility is that this is just a brute fact about our cognitive architecture. We just happen to be wired that way. This won't do, however. We can't reconstruct a notion of the *a priori* merely out of the notion of fixed cognitive architecture, because the *a priori* is a normative, not merely a descriptive notion. In order for the belief that bachelors are unmarried to count as *a priori*, it must be *irrational* ever to deny it, no matter what the evidence; and it must be the irrationality of denying it that *explains* our refusal to deny it. Mere *de facto* cognitive rigidity is not sufficient.

We have already seen one way of meeting this condition on admission to *a priori* status. In the case of beliefs true in virtue of logical form, we can explain why it is irrational ever to deny them. However, we are assuming for now that the file entry [unmarried] under [bachelor] does not engage the relevant compositional mechanisms. What we need then is an epistemic status for unrevisable statements in between true-by-virtue-of-logical-form and stubbornness.

Well, suppose we just call it "meaning postulate". That is, suppose that certain file entries are just treated, by the nature of the cognitive architecture, in a manner that endows them with the status of a meaning postulate. The cognitive system just considers it to constitute a change of meaning to give up the relevant belief (or inferential relation).

But now we are back where we started. For this really does seem to amount to just a matter of labelling, which was precisely what Quine objected to Carnap's notion of an L-rule.[24] Aside from *de*

24. Rey, in conversation, has objected to this analogy with Carnap's L-rules on the grounds that for Carnap the rules were a matter of convention whereas for him they are determined by cognitive architecture, and therefore play a significant explanatory role. I will address the question of explanatory role in section 4. As for the analogy itself, my only point concerns the normative status of the rules in question. For that purpose, I don't see how the fact that the rules are determined by cognitive architecture makes a relevant difference.

facto unrevisability and engagement by compositional mechanisms, there just doesn't seem to be anything else to distinguish the status of our file entry for [bachelor] from the corresponding entry for [water]. The Quinean argument is that *de facto* unrevisability cannot underwrite a normative notion of the analytic-*a priori*. On the other hand, engagement by compositional mechanisms collapses the analytic-*a priori* to truth in virtue of logical form, which the atomist is happy to admit.

So far I have treated the file story as an attempt to provide a purely internal criterion, and it has been found inadequate. We also have seen that a purely external criterion – grounded in the necessary relations among properties themselves – can't distinguish between the status of [water contains hydrogen] and [bachelors are unmarried]. Perhaps, then, we need to look for a mixed criterion instead, one that combines elements of the internal cognitive architecture and the external metaphysically necessary relations among properties.

The most straightforward proposal is to just conjoin the external and the internal conditions. We'll call a judgment *a priori* if it is guaranteed unrevisable (under ideal conditions) by the cognitive architecture and it reflects a necessary truth. What separates the judgment that bachelors are unmarried from the judgment that water contains hydrogen is that the former meets both conditions while the latter only meets the second condition. I don't think this proposal will work either. The principal problem is that it is *ad hoc*. The fact that our cognitive architecture is so constituted that some necessary truths are also *de facto* unrevisable does not give us a notion of the *a priori*. What we need is a principled connection between the two conditions; it must be *because* it's a necessary truth that it is unrevisable. Only if there is a connection of this sort is there sufficient normative force to justify calling a belief *a priori*. Only then do we have a case of not only unrevisable belief, but *rationally* unrevisable belief.

But how is this connection to be maintained? The necessity of bachelors being unmarried on its own can't account for it, since, as has been emphasized so often, the necessity of water containing hydrogen doesn't yield a similar unrevisability. The only way I can see that a computationalist could connect the internal unrevisability criterion with the external necessity criterion is by way of logical

form. Logic gives us an account of how features of mental states could mark necessary features of reality, by showing how syntactic relations can reflect semantic relations. In the case of [all unmarried males are unmarried] we can see how its being necessarily true plays a role in explaining our treating it so. But beyond what's guaranteed true by virtue of logical form, I don't see what else there is to which we can appeal for the purposes of such an explanation.

3.3. I think this discussion of the attempt to naturalize the *a priori*, and thereby meet Quine's challenge to a molecular theory of concepts, reveals two important lessons for any naturalistic theory of concepts. The first has to do with the relation between the so-called Quine-Duhem thesis and the possibility of an extra-logical theory of the analytic-*a priori*. The Q-D thesis is formulated various ways by F&L in their discussion, but one formulation stands out as particularly relevant here. On this reading, the Q-D thesis is the claim that confirmation relations are empirical. I take this to follow from the fact that confirmation relations largely track causal relations – smoke is evidence of fire because it's fire that normally causes it – and that causal relations are empirically determinable.

The relevance of this formulation of the Q-D thesis for us is as follows. So long as confirmation relations are determined by the causal structure of the world, and so long as we have only our theories to tell us what this structure is, then a change of theory can always engender a change in what it is rational to take as evidence for any hypothesis. Put another way, so long as the causal structure of the world is not encoded in the syntactically characterizable features of our mental representations, and so long as the computationalist is correct in claiming that it is only such features that govern mental processes, then there won't be the kind of internal criterion of rationality upon which any naturalistic theory of the *a priori* must be based. The Quinean challenge, therefore, is not merely a challenge to a semantics based on an introspective epistemology, as Rey has argued. Rather, it is a challenge to any epistemically-based semantics whatsoever.[25]

25. Excluding, of course, epistemic relations based on logico-syntactic relations.

The first lesson, then, has to do with the problem of constructing semantics out of epistemology. The second lesson has to do with the limits of constructing semantics out of metaphysics. For the naturalist-computationalist, there are two factors in play: symbols and properties. Symbols are syntactic forms – with syntactic identity conditions – and properties are objectively existing features of objects. Concepts, then, must somehow be constructed out of symbols and properties. The problem for a naturalistic molecular theory of concepts is that the naturalist doesn't acknowledge either a realm of concepts distinct from (or intermediate between) symbols and properties, or a form of privileged epistemic access to the realm of properties itself. Thus *conceptual* truths – as opposed to *necessary* truths on the one hand, and *logical* truths on the other – have no clear place to reside within the naturalist's framework.

4. If I'm right that prospects for a molecular theory of concepts are not good, and if F&L are right in rejecting holism, then we'd better take a closer look at atomism. It seems to me that we are not only led to consider atomism by process of elimination, but that the lessons just drawn from the discussion above should lead naturalists to consider atomism in a genuinely favorable light.

Rey (forthcoming) presents two sort of objections to atomism. First, he claims that the position is just implausible on the face of it. How could it even be possible for a subject to have a concept like DEMOCRACY without having any other concepts at all? Second, he claims that atomistic concepts serve no explanatory role. Appeal to molecular concepts can contribute significantly to explanations of important psychological regularities; explanations we lose if we type concepts atomistically.

In their discussion, F&L attempt to dispel some of these worries, but they don't really address the question directly; especially the second worry about explanatory import. So what I propose is to expand on their remarks, and, putting them together with the lessons of section 3, draw a picture of the atomistic theory that is, if not compelling, at least plausible.

There are three distinctions that bear on this question: 1) between intentionality and mentality, 2) between metaphysically necessary (or constitutive) conditions and nomically necessary conditions, and

3) between metaphysically necessary connections among properties and analytically necessary connections among concepts.[26]

The first distinction helps to undermine the intuitive resistance to atomism by granting that in order to have full-fledged mental states, or even mental states of certain types (e.g. beliefs and desires), it may be necessary that a subject meet rather stronger conditions than the atomist requires for the possession of a concept. Perhaps minds are necessarily reasoning systems with preference functions, or something of the sort, and of course this demands a lot of conceptual/representational machinery. Thus a subject that could entertain only one concept might not be an appropriate subject for the attribution of beliefs and desires.

But it is one thing to claim that in order to be a mind one must meet these stronger conditions, and it is another to claim that in order to have intentional, or semantic content such strong conditions must be met. It is the latter claim that the atomist denies, and, F&L emphasize, all that she need deny. Of course one might wonder then about the relevance of the atomistic notion of intentional content to the individuation of mental states. The point is this. While it may take stronger, non-atomistic conditions to count as a *mental* state, these conditions are not constitutive of the *intentional content* of mental states. It is fallacious to infer from the fact that x can't be an F without being a G, that being a G is constitutive of every one of x's other properties.

The second and third distinctions work together to undermine the intuitive resistance to atomism, in a complementary way. The atomist employs them as follows. Sure, she argues, it seems wild to claim that someone could have certain concepts without having a host of its inferential liaisons (e.g. CAT and ANIMAL, or BACHELOR and UNMARRIED, or FORCE and MASS). However, there are two crucial points to keep in mind here: First, though the atomist is committed to the logical possibility of (say) having the *concept* BACHELOR without the *concept* UNMARRIED, she is not committed to the logical possibility of an *unmarried bachelor*. As emphasized above, denying analytic connections between concepts

26. All three distinctions are employed by F&L, though in different degrees of explicitness.

in no way entails the denial of necessary connections between properties. Thus, some of the implausibility that attaches to the claim that it is possible to have certain concepts without certain others might be diagnosed as displaced reaction to the claim that it is possible for the relevant property to be instantiated without its necessary concomitant.

Secondly, though (again) the atomist is committed to the *logical* possibility of having the concept BACHELOR without the concept UNMARRIED, she is not committed to its *nomic* possibility. That is, depending on the sort of relation that must obtain between the symbolic representation of a concept and its satisfaction conditions in order for the latter to count as the semantic content of the former, it may not, as a matter of nomic constraint, be possible for someone to have certain concepts without their inferential liaisons. The reason this may not be possible is that possession of the inferential liaisons may be part of the mechanism which sustains the relevant semantic relation. For example, if you don't believe that bachelors aren't married, you may not be (nomically) able to stand in the relation to bachelorhood that is necessary for your term [bachelor] to express the concept BACHELOR. The atomist's point is that this is a nomic constraint, not a logical/metaphysical constraint, and therefore the concept UNMARRIED is not thereby metaphysically constitutive of the concept BACHELOR.[27]

One way to see the force of this atomist move is to consider a fairly natural situation in which the allegedly constitutive conceptual constraints and the nomic constraints come apart. Many of the examples that seem most difficult for the atomist involve concepts which are embedded in theories; concepts of properties to which we only have theoretical access. So one might question whether one could have the concept of a pi-meson without some knowledge of physics.[28] Of course the atomist attempts to handle the intuition that

27. This argument is just the other side of the coin with respect to the objection to Devitt in section 2.2. He wanted to make these nomically necessary conditions constitutive of content. The point here is that they do all the work they need to – in particular, accounting for our strong intuition that you can't have the concept of X without the concept of Y – without playing a constitutive role.

28. This example occurred in an earlier draft of Rey's paper.

no one could have that concept without knowledge of physics by emphasizing that only through such knowledge could the subject stand in the appropriate relation to pi-mesons.

But notice that we do have sensory access to all sorts of properties that are themselves highly complex in nature, and the full understanding of which requires quite a lot of theoretical sophistication. Take colors, for an obvious example. Supposing for the moment an objectivist position on the metaphysical status of colors, the color of an object is the reflectance function of its surface.[29] Normally sighted subjects can pick out the same reflectance function reliably without knowing anything about the theory of color vision or light. Of course they'll call it "pink" (or [pink]), but the property itself is a reflectance function.[30]

Does the theoretically uninitiated share a concept of PINK with the color vision theorist? For that matter, if we had a sense organ sensitive to pi-mesons the way our retina is sensitive to photons, would we thereby automatically have a concept of a pi-meson? The atomist and the molecularist will disagree here, and I don't think the question can be settled on general intuitive grounds. But that's my point. Atomism isn't obviously crazy, and whether it is ultimately tenable depends on the theoretical employment of the concept of a concept; which brings us precisely to the second issue – the worry about explanatory import.

Again, to fix ideas, let us suppose that the atomist in question is one who holds that the semantic content of a mental symbol (predicate) is a property, the identity of which is determined by the property's standing in a specified nomic relation to tokenings of the symbol. So the content of [cat] is the property of being a cat in virtue of there being a nomic relation of the appropriate sort, call it 'R', between tokenings of [cat] and that property. On this view, then, a subject counts as thinking there's a cat in front of her when she

29. See Hardin (1988) for arguments in favor of subjectivism about color, and Hilbert (1987) for arguments in favor of objectivism.

30. Here we see another echo of our earlier discussion. Objectively speaking, pinkness and the reflectance function are the same property. What distinguishes the casual color observer from the vision scientist is the structure of their representations of this property.

stands in the "thinking-that" computational relation to a mental sentence like [there's a cat in front of me].

Now what is supposed to be the problem about explanatory import anyway? The idea seems to be that we should be able to derive psychologically significant consequences from the attribution of a mental state, and this doesn't seem possible on the atomist account. If we know that to have the concept CAT is to subscribe to a little mini-theory of cats, as it were, then from the fact that Jones thinks she is currently confronted by a cat I might be in a position to infer various consequences – e.g. that she thinks there's an animal in front of her, that she thinks that what's in front of her had to be born, breathes, and a host of other consequences. But if all I know is that she is now tokening a symbol that bears R to cathood, none of these consequences follow, and it seems I won't be able to explain any of her actions or subsequent mental states by reference to this mental state.

It seems to me, however, that molecular theories have no real explanatory advantage over atomistic theories. To see this, we need only appreciate again the distinction just drawn between inter-symbolic relations that are constitutive of a concept and those that are nomically necessary for expressing it. My point is that all the explanatory work we need done can be done, and in fact is better done, by appeal to the latter sort of connection.

For example, suppose we see that Smith is looking straight ahead at a cat. We surmise that she believes there's a cat in front of her. What can we now predict about her other beliefs and behavior? Well, certain predictions would depend on further information on anyone's story. So, for instance, whether Smith is going to pick up the cat and start petting it or immediately turn around and head in the other direction would depend on such facts as whether Smith was a cat-lover, on the one hand, or perhaps had a cat-allergy (or, to be precise, *believed* she had a cat-allergy), on the other. Presumably the molecularist is not going to treat these collateral beliefs/attitudes as constitutive of Smith's CAT concept. (To do so would clearly involve a capitulation to holism. Or, to be more precise, to treat *all* of these collateral attitudes as constitutive of her CAT concept would be to capitulate to holism.)

Well, what about other possible predictions, such as whether or

not Smith believes that what's in front of her is animate? Suppose the molecularist treats the belief that cats are animate as constitutive of Smith's CAT concept. Then, of course, it will follow from Smith's believing there's a cat in front of her that she believes that what's in front of her is animate. For the atomist, on the other hand, this inference doesn't hold. So, it looks as if there's a prediction (or, an explanation) available to the molecularist that isn't available to the atomist.

But, we have to be careful here. How do we know that Smith believes there's a cat in front of her in the first place? For the molecularist, as opposed to the atomist, before we can justify that ascription, we first have to know that Smith believes that the sort of thing she sees in front of her is animate. In other words, on the molecular view we can't predict from the information that she is tokening [cat there] in response to seeing a cat that she also thinks [animate object there]. Once we know she's already tokening [animate object there], then we can say her [cat] token expresses CAT. But by that point, even on the atomist's story, we *already have* the information that she believes what's in front of her is animate, and thus there is no explanatory advantage for the molecularist.

It isn't just a matter of drawing psychological consequences from observational conditions. The point is quite general. Suppose we have a psychological law of the form, if x believes P then she'll believe Q, and the molecularist claims that believing R is a conceptually necessary condition for believing P, while the atomist denies this. Let's also assume that believing R is nomically necessary for believing Q. So, the form of the law will be simpler for the molecularist, since she needn't mention x's believing R explicitly, whereas the atomist might have to add the condition that x believes R as well as P. (I say "might" because it may be nomically necessary that x believe R if she believes P, in which case it could go without saying.) But when it comes to applying the law, both the molecularist and the atomist will be in the same position; namely, both must check x for both the belief that P and the belief that R before inferring that she believes that Q.

The molecularist might respond by arguing that there is still a difference in explanatory insight, because she can explain why anyone who believes P also believes R. The reason is that believing

R is partly constitutive of believing P. But what extra explanatory work does this really do? For now we want to know why people who meet the atomist conditions for believing P also tend to meet the molecularist conditions for believing P, and *this* is going to have to be a matter of empirical law, since presumably it's an empirical issue whether anyone meets the latter conditions. But then the atomist can just appeal to those very same laws to explain why people who believe P also believe R.

The moral of the story is this. On the computational model, the mind is a system of symbols – syntactically individuated – that causally interact with each other and with the external world. These are the basic facts with which any account of intentional content based on the computational model must start. Given the availability of this web of causal interactions and dependencies, it's not clear what explanatory or predictive advantage could accrue to the molecularist merely by virtue of designating a subset of the internal dependencies as conceptual ones.

One final note on the question of explanatory import. So far I have tried to show that the molecularist derives no explanatory advantage from her view. I have not said anything about the positive role that intentional content plays for the atomist. In fact, one plausible reaction to this discussion is to say that it demonstrates the explanatory futility of an intentional taxonomy of psychological states altogether. Rather, we should stick to a syntactic taxonomy, relegating intentionality to the realm of the pragmatic and unscientific.[31]

I think this is an important challenge, and not one that I can take up here.[32] My only claim is that if one wants to make an intentional psychology work, an atomistic theory of intentional content should not be ruled out of court too quickly. It has the resources to handle what look initially to be insurmountable problems, and, what's more, it fits more naturally than a molecular theory into the scientific naturalist's overall metaphysical and epistemological framework. It

31. This is essentially Stich's position in Stich (1983).

32. In Fodor's recent unpublished work, he tries to meet this challenge. Also see Antony (1990), who claims that we need content to account for behavioral success.

could do a lot worse.[33]

33. I'd like to thank Louise Antony, David Auerbach, Michael Devitt, Jerry Fodor, Georges Rey, and the members of the Triangle Unethical Society for helpful discussion and comments on earlier drafts. I also want to acknowledge my debt to NEH, and the members and directors of the 1992 NEH Summer Seminar on Meaning Holism at Rutgers University.