A Dilemma For Asymmetric Dependence

Joseph Mendola

Accounts of mental content rooted in asymmetric dependence hold, crudely speaking, that the content of a mental representation is the cause of that representation on which all its other causes depend.[1] To speak somewhat less crudely, such accounts, hereafter "AD accounts", hold 1) that (a) any nomic relations which obtain between, on the one hand, a mental representation and, on the other hand, various causes and possible causes of that representation which are not its content, are dependent on (b) a nomic relation between that representation and its content, but 2) that the latter nomic relation is not dependent on any of the former relations. This paper argues that AD accounts are false.

There is more than one AD account; there are a range of possible accounts which share that general form. Yet all of those alternatives are false. The problem is a dilemma defined by reference to three axes of difference among possible AD accounts. Any position along the first axis requires a problematic position along one of the other two. First, AD accounts may deploy nomic relations which differ in scope, which apply for instance to all individuals or merely to particular individuals and hence have in that sense relatively wide or narrow scope. Second, AD accounts may deploy different conceptions of the nature of the syntax of mental representations. Third, AD accounts may deploy different accounts of the grounds of the nomic relations which bind that syntax to target contents. I will argue that AD accounts which deploy wide scope laws cannot fit with a reasonably plausible account of the nature of the syntax of mental representations, and that AD accounts which deploy narrow scope laws cannot fit with a reasonably plausible account of the grounds of the nomic relations which bind that syntax to target contents.[2]

Section 1 is a brief exposition of AD accounts and of their motivation. It also sketches

some of the axes on which they may differ, and in particular the three crucial axes just noted. Sections 2 and 3 reveal problems along the other two critical axes faced by AD accounts which deploy, respectively, wide or narrow scope laws.

1

The basic idea of information-based accounts of mental content, of which AD accounts are a variant, is perhaps best conveyed by a simple caricature. That is the view that X represents Y, that X means Y or has Y as content, just in case Y is the cause of X. For instance, an internal brain state of a certain neurophysiological type might be held to represent cows if it is caused by cows.

This is a simple caricature in various ways. For instance, the causal relationship which X and Y are supposed to bear may be somewhat more refined.[3] But the primary axis on which information-based accounts are customarily distinguished maps their responses to a certain characteristic objection to their crude form, a general objection which comes in a variety of specific types, some of which are called “error” or “disjunction” problems.[4] The general problem is this: Clearly some causes of any representing state are not its content. For instance, if there is some type of brain state, say part of some perceptual system, which represents cows in the vicinity, it may also be caused by horses on dark nights. Plausible defense of informational accounts requires more closely specifying the relevant causes. For instance, we might claim that the content of the representing state is the cause which it had during a specific learning period, or which it is its evolutionary telos to be caused by. But the central feature of AD accounts is a different conception of relevant causes, a conception whose primary champion is Jerry Fodor.[5]

There are forms of AD accounts which incorporate a concern with actual causal histories, as suggested by my initial crude characterization, and to which we will return. But the paradigmatic forms rely predominately on nomic relations, which may be merely hypothetical in the sense that they link merely possible causes to merely possible effects.[6] The basic idea of such

AD accounts is that the target contents of mental representations of type R are those things of type Y, among all the things which are nomically related to Rs, such that all the other things are nomically related to Rs because Ys are, and such that it is not the case that Ys are nomically related to Rs because any of those other things are. In other words, Ys are the causes or possible causes on which the others are dependent, and which are themselves not dependent on the others. The other causes are hence said to be asymmetrically dependent on Ys, in the sense that if Ys had not been causes of Rs, then the other things would not have been, but the opposite dependency doesn't hold. Such a nomic relation guarantees that Ys would cause Rs in various hypothetical situations. So it is true that, counterfactually, if a Y occurred in certain circumstances, then an R would occur. But this nomic relation also supports counterfactual conditionals of another sort, such that if Ys had not been nomically related to Rs, then the other causes would not have been. To apply this general strategy to our case: The basic idea is that if mental representations of cows are caused by horses then that is because cows also cause or would cause them, but not vice versa.

The primary motivation of AD accounts is to get the cases intuitively right, and hence block the primary objection to crude information-based accounts. While Fodor, in his guise as the primary partisan of AD, has given other positive arguments for AD, there are reasons for caution about their significance. First, they involve public language cases, while AD accounts are primarily accounts of mental content: (Fodor 1987, 107-108) says that if one mistakenly says "horse" of cows it is only because it means horse, because one says it primarily of horses. And (Fodor 1990, 96-99) claims that the practice of paging people depends asymmetrically on the practice of naming them. Second, an unsympathetic reading of these passages would suggest that Fodor's intuition that the adventitious causes of something are dependent on the content-relevant causes is merely his intuition that the first are from the semantic standpoint accidental, in other

words that they are causes which aren't contents.[7] This may make us suspicious that such an account cannot really be developed into a reductive account of content, but working out this suspicion will require that we look at the details, and see that AD accounts cannot get the cases right either, at least without deploying some very implausible metaphysics.

There are a number of ambiguities or generalities in the sketch of AD accounts I have given. If resolved to clear specificity in various ways they reveal a range of accounts which might properly be considered versions of AD, which differ on various axes.[8]

A plausible AD account must occupy a particular location on the first axis. In his development of AD, Fodor uses "nomic" in a stronger sense than any I've yet explicitly introduced, such that nomic relations are not only causal relations which imply counterfactuals but also must link properly projectible properties, elite properties which in some sense "cut nature at its joints". This matters. There are a variety of apparent counterexamples to AD as so far specified, which suggest that the target nomic relation, relating a mental representation and its intuitive content, is asymmetrically dependent on other relations involving parts of that content or things more general than that content or things intervening between that content and its representation.[9] We will discuss a central instance of this objection in section 3. So AD as so far explicitly characterized fails to yield intuitive contents for thoughts. Fodor's crucial response has been that such alternate properties aren't proper candidates for nomic relations even if they seem to suggest the relevant counterfactuals, that they aren't properly projectible and elite.[10] On the basis of passages like these, it might be more accurate to say that AD uses two resources together to block the characteristic problem of information-based accounts of content, a restriction to nomically elite properties and also asymmetric dependence itself, and not just the eponymous one.[11] A plausible AD must adopt this fix.

A second axis on which AD accounts might conceivably differ but on which plausible

accounts must occupy a particular point is reflected in Fodor's reasonable insistence that the counterfactual-supporting asymmetric dependencies between nomic relations which are constitutive of content are synchronic rather than diachronic. They are dependencies which obtain at the time when the content is possessed by the mental representation and not dependencies which obtain over time. Such a clause is clearly helpful in disposing of counterexamples, because, for instance, a mental representation's having a content may be diachronically dependent on some history of education which is intuitively not its content. But it is perhaps worth noting that it may make AD accounts a somewhat odd fit with popular historical theories of reference.[12]

There are three other axes of possible AD theories which will structure our discussion. They are relevant to determining the exact truth-conditions of various claims about asymmetric dependence which AD accounts must make.

First, it is sometimes ambiguous in at least Fodor's formulations of AD exactly to whom particular context-fixing nomic relations are supposed to apply. Do individuals instantiate different content-determining laws, even perhaps different laws in different moments of their individual lives? Or rather do the relevant laws apply across a linguistic community, or a species, or even to all cognizers? Whatever Fodor's commitments,[13] a consistent form of AD must adopt one such account. I'll call this "the scope issue". Laws which apply to all cognizers or across a species have relatively wide scope, while those which apply to individuals or to individuals at moments have relatively narrow scope.

The second structuring axis is spanned by various accounts of the nature of the syntax of mental representations, the identity of the mental representations to which an AD account is to ascribe content and hence which enter into the relevant nomic relations with target contents. It is again sometimes not sufficiently clear in Fodor's various versions of AD what fixes the identity

of the mental representations to which content accrues,[14] and again, whatever Fodor's commitments, a consistent form of AD must adopt one such account. Call this "the syntax issue".

The final structuring axis involves a quite general issue, regarding exactly what grounds the truth of the relevant nomic claims, and of the claims of dependency among those nomic relations, which an AD account deploys. I'll call this "the grounds issue". Not only are there a spectrum of possible views on this issue, there are many different specific issues about such truth grounds we might discriminate, and which suggest a variety of different sub-axes to organize discussion of AD accounts. For instance, there are quite broad worries about nomic relations in general and also quite specific ones about those relevant to content in particular. For now, I will merely note three issues of this general sort which will be significant later on.

As we have seen, it is important that the relevant nomic relations are relations among properties of a rather elite sort, those which cut nature at its joints and are hence projectible. But even this leaves space for a number of relevantly alternative positions, depending for instance on if and exactly how abstract entities like transcendental universals are deployed to underwrite nomic relations between elite properties, to single out certain properties as elite. This is the first grounds issue to which we will return.

Second, Fodor works within the externalist tradition in at least recent writings, and the general dialectical position of AD as a development of externalist information-based accounts of content suggests that this is a plausible home for AD. But there are large elements of AD which might be made consistent with an internalist conception of what grounds content-fixing nomic dependencies. It is possible to field what Fodor calls a "pure informational account", which requires no actual causation at all but merely nomic relations which ensure that causation would hypothetically occur under various conditions. This in one sense is the most characteristic form of AD account, and something like this is suggested by some of Fodor's own remarks, notably about

unicorns. See (Fodor 1990, 100-101). But, on the other hand, Fodor sometimes deploys formulations which require, if some mental representation is to represent X, that some of its type be actually caused by Xs. See (Fodor 1990, 121). And he also characteristically requires, for instance at (Fodor 1990, 118), that some of its type *not* be caused by X's, that they have other causes, for it to mean X. This second restriction is motivated by the thought that content must accrue to mental representations robustly, that even if caused by an inference or perceptual error a mental representation must retain its content, and it has also been deployed in response to counterexamples. It is such requirements about instanced or uninstanced causal relations which most obviously would make a version of AD externalist, so that conditions external to the skin of someone are part of what constitutes that they think as they do. And there is more than one detailed way that externalist resources may be deployed by an AD account, not only to fix the nature of actual causal histories in the manner already noted, but perhaps also to fix the content of hypothetical causal relations. For instance, if there is only H-O-H around in the environment and no XYZ, then it may seem that even if one hasn't actually interacted with H-O-H, still in some relevant sense one could have only so interacted.

Third, and most important, physicalism is one constraint of plausibility which governs the nature of relevant nomic dependency claims. Not all AD accounts are physicalist, but plausible AD accounts are. And in turn physicalism generates further serious constraints on AD. The implications of physicalism are the third grounds issue we will trace.

Let me mention some problems physicalism may pose for AD accounts which will be relevant later, especially in conjunction with our first grounds issue. As we noted, what are in some sense characteristic counterexamples faced by AD require that only genuine, full-blown nomic relations among certain elite properties, though perhaps including uninstantiated properties like that of being a unicorn, are relevant to content. See (Fodor 1990, 93-95 and 100-102). One

might, I think, legitimately worry that if we take this talk of at least uninstantiated elite properties seriously, it is in some immediate tension with physicalism as an ontological doctrine. And even if, at least for instanced properties, we might reasonably hope that a distinction between properly projectible properties on the one hand and not properly projectible but yet counterfactual-supporting properties on the other might be made consistent with physicalism without deploying transcendental universals, still we might also reasonably worry that the peculiar eliteness of fully nomic properties in at least Fodor's formulations surreptitiously involves indeed even more than mere projectibility, perhaps indeed something so rich that it cannot be delivered even by those forms of realism about universals which can arguably and at a stretch be made consistent with physicalism. For now this is merely a suspicion, but we will return to this issue.

There is a sense in which all these axes of possibility for AD accounts represent a richness in its basic conception. But we will see as we proceed that all the specific alternatives with which we are left after we resolve our three principal structural axes are problematic. We will see that resolving the scope issue towards the general end of the spectrum, say in favor of species-wide laws, makes it impossible to resolve the syntax issue in a plausible way which doesn't generate debilitating problems for AD accounts. And we will see that resolving the scope issue towards the specific and individual end of the spectrum generates corresponding difficulties in at all plausibly resolving the grounds issues. That is our dilemma for AD accounts.[15] The first horn is the primary concern of the next section; the second horn is the primary concern of section 3.

A more detailed outline of the argument may be useful. I presume that if an AD account is true, then it assigns intuitive contents to mental representations. And AD accounts must deploy either wide scope or narrow scope laws. But we will see in section 2 that if the laws are wide scope, then mental representations are typed by causal-functional roles, and that if mental representations are so typed, then AD accounts assign them unintuitive contents. And we will see

in section 3 that if the laws are narrow scope, then AD accounts also assign unintuitive contents, at least if nomic dependencies have plausible grounds.

2

Fodor holds that psychological states with content involve inner mental representations in a language of thought often called “Mentalese”, a language with determinate and combinatorial syntactic structure. There are certain token words of certain specific types, as for instance “dog” and “dog” are two instances of the same type of word, which combine together to form sentences with specific sorts of syntactic grammar. And any AD account is committed to some sort of typing of the mental representations, of the vehicles of representation. I’ll call that the syntax of the language of thought.

The nature of that syntax is crucial for AD accounts, because it is to words or sentences of the so-called language of thought that content accrues, and which are hence one set of the crucial relata bound by the nomic relations which fix content. What is supposed to determine the syntax of someone’s mental representations? Fodor himself has long and consistently held that syntax is fixed by internalist resources, and that it is a “formal” resource, which is to say that it is fixed independently of semantic considerations. Whatever Fodor’s commitments, these seem to be reasonable constraints on the nature of mental syntax, which we should accept. But there are three more specific accounts of syntax which we should consider. It is notable that all three are suggested at some point in Fodor’s writings, and that this gliding by the principle partisan of AD hides some of the difficulties we will shortly trace.

Written words in natural languages are typed more or less by shape, spoken words by sound profile. Such a very simple and concrete account of the syntax of the language of thought is also possible, and indeed it is suggested, though usually at arm’s length, by some things Fodor

says. (Fodor 1980, 64) claims that, at least intuitively and metaphorically, “formal operations apply in terms of the, as it were, “shapes” of the objects in their domains.” (Fodor 1990, 112) says that for “the purposes of a naturalist semantics the only nonquestion-begging reading of “cow” is #c^o^w# . . .” where “cow” is a mental representation and #c^o^w# captures an “orthographical/phonological sequence”. That is the first account of syntax we should consider, that it is typed by shape or something equally concrete.

But it is somewhat more plausible, and indeed sometimes Fodor seems to propose, that syntax is typed by an abstraction from concrete shape. “The syntax of a symbol is one of its higher-order physical properties. To a metaphorical first approximation, we can think of the syntactic structure of a symbol as an abstract feature of its shape.” (Fodor 1987, 18).

And yet even in a footnote to that passage there is a still more plausible third suggestion, that “[a]ny nomic property of symbol tokens, however--any property in virtue of the possession of which they satisfy causal laws--would, in principle, do just as well. (So, for example, syntactic structures could be realized by relations among electromagnetic states rather than relations among shapes. As, indeed, it is in real computers.) This is the point of the functionalist doctrine that, in principle, you can make a mind out of almost anything.” (Fodor 1987, 156). This at least points towards typing syntax by causal-functional role.[16]

Whatever the niceties of Fodor exegesis, such a range of accounts of mental syntax is possible, and we need to consider the form of asymmetric dependency in light of each. As I’ve said, the nature of syntactic typing and the proper form of an AD account are linked, since a syntactic item is one of the two relata bound by a relevant nomic relation. I will argue in this section that more abstract and relational accounts of syntax are much more plausible if the language of thought and hence relevant content-fixing laws are supposed to be invariant among, say, all humans, that is to say if such laws have wide scope.[17] But I will also argue here that this

generates a serious problem when conjoined with AD as an account of content. Problems for more narrow scope accounts will be most evident in the next section, when we discuss the other relata bound by AD.

Let's begin our discussion with the simplest case, syntax-typing by literal shape, or by something equally concrete like a specific sort of quite concretely-specified neural firing pattern. My main point about this is that it is extremely implausible that the same shapes or firing patterns are present across all cognizers for the representation of, say, water, or even of a particular taste. This is assured by the plausibility of the intuitions which underlie the traditional multiple-realization problem. And such a very concrete type-theory certainly would not be supported in this strong and implausible form by the familiar and widely accepted arguments for the language of thought thesis: arguments from productivity, systematicity, and from the nature of mental processes. (Fodor 1987, 135-154).

Nor is the language of thought thesis very plausible if committed to such a fixed concrete typing even merely across a single species. Fodor himself suggests that there are primitive representations not just for sensory properties but for properties like doorknob and firehydrant. (Fodor 1998a, 89-145). But it simply isn't plausible that all humans are fitted out with a mechanism which will assure exactly the same firing pattern or shape is involved in the mental representation of such things. Even if, as Fodor thinks, we are innately fixed to generalize from stereotypical doorknobs to possession of the doorknob concept, we aren't fixed to represent them in exactly the same way, by exactly the same shaped tokens. How could we be? Even if the form of the syntax is fixed by the stereotypes, we interact with somewhat different stereotypes. Perhaps it might be claimed that we somehow generalize from different stereotypes to a single concept, and then that determines the shape of the relevant representation, but that would violate the formality condition.

Of course, AD accounts can abandon Fodor's other commitments. Perhaps there are only primitive representations for sensory properties. Still, brains, even brains of identical twins, plausibly exhibit different patterns of excitation and different shapes at the suitably fine-grained level. And we shouldn't forget the traditional functionalist worry that, absent the patterns of connectivity which surround it, a particular concrete firing-pattern would plausibly retain no content at all, whatever its asymmetric dependencies.[18]

We can reasonably conclude that if the laws which fix the content of a piece of syntax lock content to concrete shape or particular firing patterns, then the laws are laws for individuals, or even perhaps merely for individuals in brief periods of their lives. This, we will see in the next section, creates certain problems for AD accounts, but not those which immediately concern us.

So let's consider the next possibility, which is syntactic typing by something a bit more abstract than literal shape. It is likely that there are two types of account of this kind, those which deploy causal-functional relations to other pieces of syntax, and those which instead type each sort of syntactic token intrinsically. So it is likely that this middle type reduces to our other two types, with relevant difficulties intact. Still, because it is perhaps somewhat more plausible to claim that some intrinsic abstraction of shape is common to a bit of syntax shared by all cognizers or at least all humans than that shape itself, let me consider such a view directly.

The question of course is which abstractions are relevant. While abstract shapes can be realized in a variety of ways, still it may seem that the disjunction of realizers should not be "wild" or "open". The fan of AD might say that those abstract shapes which enter into laws are the relevant ones, are properly nomic. But there are two ways to understand this.

On one hand, perhaps it is supposed to be the case that whatever types of syntax there are involving abstract shapes is set by the facts about whatever it is that is nomically tied in the appropriate way to represented properties, that we find the syntax by starting with what is

represented and working back along the content-determining nomic relations it enters into. But this violates the very plausible formality condition, which is that syntax is fixed by other than semantic resources. If “cow” is just whatever represents cows, then Fodor’s project and most of our current reductive attempts at mental semantics fail outright.

On the other hand, perhaps the laws which fix syntax are independent of those which fix content. There are various ways this might be so. As Fodor himself suggests in response to Kripke’s worries that actual dispositions cannot fix causal-functional roles, it might be that certain types of syntax are fixed by the nature of the projectible and elite properties there are in the world, properties which are candidates to enter into content-determining laws. But of course to the degree that the relevant syntactic typings are not invariant across all cognizers, these elite properties would need to be fairly locally-based and locally-applicable. Narrow scope projectibility, relevant to a particular individual psychology, seems an inappropriate resource to constitute eliteness if it requires violation of the formality condition, and on the other hand some sort of projectibility in that individual psychology supporting laws defining syntactic identity independent of content would apparently simply collapse this alternative into a narrow scope version of the causal-functional form of syntactic-typing which we are soon to discuss. Relevant syntax might alternatively be fixed by specific and properly individualized Platonic forms hovering over particular cognizers or specific irreducible Aristotelian teli pushing out from within them, but this seems quite like dualism by another name. In any case, as we will see in section 3, narrow scope accounts face debilitating problems.

So perhaps we should consider relatively wide scope syntax of this general sort, and deploy more physicalistically-acceptable forms of platonism or a wider sort of projectibility to get the elite properties it requires. But the problem with these routes is that it is implausible that a particular piece of syntax in all sorts of cognizers has an at all reasonably concrete material

specification anyway, for reasons we have already considered. If that is required by AD, then it seems plausibly false. The only form of plausible wide scope laws of this sort are abstract enough to count as full-blown causal-functional typing.

The most plausible forms of our second sort of syntactic-typing have collapsed into the third. And we certainly seem forced to that last and relational alternative for typing syntax if we are to avoid very narrow scope laws. So let's focus there. Let's presume that a piece of syntax is what it is because it plays a certain causal-functional role.[19] Perhaps paradigmatically, this role mirrors its narrowly logical role, its role in proper formal deductive inferences.

This causal-functional conception may seem to threaten a vicious circularity, but no more does it do so than any causal-functional scheme which characterizes some level of entities all at once, by their role in the overall system. There are of course a variety of different forms this causal-functional typing might take. First of all, the typing might be unique for individuals or for moments of individual lives, or correspondingly more social or universal; it might be narrow or wide in scope.[20] If the laws relevant to content are individual, of course, we will face the problem promised in the next section. Perhaps, however, it might be that one could abstractly characterize a constituting causal-functional role for a piece of syntax which is constant across all individuals, at least say of a species. So, at least for all we know so far, reasonably wide scope laws are plausible and consistent with AD. They just require causal-functional syntactic typing.

The primary argument of this section, however, is that plausibly wide scope and hence causal-functional syntactic typing creates a debilitating problem when conjoined with an AD account of content. Causal-functional typing and asymmetric dependence together yield unintuitive content.

By an AD account of content, "cow" means cow just when i) it is nomically related to cow, and ii) if it is nomically related to anything else then that other relation is asymmetrically

dependent on the first. But “cow” can’t mean cow unless “cow” exists, at least in whatever potential way any currently uninstanced word in the language of thought exists. And if “cow” is to so exist, on the conception of syntactic-typing which we are currently exploring, then it must be such as to be nomically related to whatever is specified in the role which fixes its syntactic type. And in addition it would be a “cow” (given the plausible formality condition) even independent of semantic considerations, even if there were no cows and it were not nomically related to cow in the way that it is. “Cow” cannot be nomically related to cow unless “cow” exists in the relevant sense, and it cannot so exist unless it is nomically related to its typical syntactic cause. And by the formality condition it can and would be nomically related to that syntactic cause even if it bore no nomic relation to cow. So “cow” in fact means typical syntactic cause of “cow” and not cow, according to AD accounts.

This objection is pressing whether the laws fixing content apply to individuals, species, or across all cognizers, as long as syntax is typed in a causal-functional way. But we have seen that such a form of typing is plausibly required if wide scope nomic relations fix mental content. Or at least that is so absent enormously implausible claims about constancy of the shape of particular mental representations across cognizers.

How might the partisan of AD reply? One source of possible objection to my argument is suggested by an analogy between my case and others which have been discussed in the literature. Those other cases do not seem debilitating, so neither may my argument. But the analogous cases are not couched in such a way as to evade defensive deployment of elite properties. For instance, while the robustness clause of the standard Fodorian formulation of AD assures that mental representations will have non-content causes, still arguably there is no law linking such causes to the representation. Seager has argued that at least when AD is applied to natural languages malapropisms provide counterexamples somewhat analogous to my own, and Adams and Aizawa

have stressed adventitious causes as problematic for AD accounts. But against these analogous cases, a partisan of AD can deploy the reply that the alternative property which AD threatens to nominate as content is not nomic. Still, in the case we are considering it must be nomic, since the syntactic types must be nomic according to AD to be bearers of content. So the characteristic response to the analogous cases can't work against my case.[2][1]

Another possible source of objection is a concern that if my argument were successful, it would also be successful against functionalism generally, and we have lots of reasons to believe in functionalism. Functionalism allows that various causal-functional roles accrue all at once to particular bits which instance them, so why can't the nomic laws which create syntax and those which fix content accrue all at once to some brain state? But notice that the formality condition assures that the various laws in question in our case are not all on one level in this sense, and cannot accrue all at once. Those relevant to syntactic identity come first and independently. That additional dependency structure is a difference between the situation which AD faces and traditional functionalist conceptions, and it is also what generates the false assignment of content according to AD.

Despite the fact that to abandon the formality condition is to give up on one of the central working assumptions of current discussions of content, still one might wonder in this context why we shouldn't give it up, or at least interpret it less rigidly than I have. Alternatively, one might attempt to modify AD to rule the particular dependency which generates the problem at issue irrelevant to content. So let's consider such responses to our argument, which require modification of the basic account.

It may help to think of what must be presumed by AD accounts in a stratified way, such that there is an in some sense staged process: There is a weak causal-functional role which fixes syntactic identity, say across all humans. That's the first part of the first level. Then, in the second

part of the first level, we somehow causal-functionally characterize the particular role in which such a piece of syntax will be primarily caused in the content-fixing dependencies, say as an element of the belief-box. So now we have a strong causal-functional characterization of one of the relata of a relevant nomic dependency, and can move on to the second stage. In that second stage, the instancing of a particular concrete piece of syntax of some sort, fitting the necessary weak role, is caused in a particular causal-functionally characterized place, say in the belief box, and hence fitting its strong role also, by the target content, and that is a matter of law, and indeed of a law on which other related nomic relations asymmetrically depend.

Note that it must be that the basic law links the target property qua that property on the one hand with the Mentalese word qua meeting the necessary strong causal-functional characterization on the other. Now suppose we modify AD accounts so that the content is the second-stage cause on which all other second-stage causes depend, the cause on which all the other causes depend asymmetrically once the identity of the second-stage causal relata are fixed, partly of course by the causal-functional role which fixes relevant syntactic identity. What's wrong with that?

Let me stress that this is a modification of extant AD accounts. But it also makes claims about the forms of nomic relations and their dependency which are implausible even if coherent, and in any case a very bad fit with physicalism. If AD can escape one horn of our dilemma in this way, it runs more quickly than expected onto the second, which involves the grounds issue. Though the grounds issue is our focus in the next section, let me briefly explain.

AD is plausibly constrained by physicalism, and yet at once deploys special science laws to fix content and claims that certain properties cannot figure into those laws because they fail to be projectible. Not everyone accepts that this is a coherent package, but since some disagree, I waive those worries here.[2][2] As we will see in the next section, it is also important to AD that the

laws which govern the correlation of mental representations and contents at least across individuals be indifferent to individual differences of belief and perception, in a way which may seem largely miraculous, but is still roughly comprehensible as consistent with physicalism since there can intuitively be macrocausal relations which are indifferent to variations in the individual detail of the mechanisms which instantiate them. But there is greater tension with physicalism in the case at hand than mere multiple realization of familiar sorts suggests. One of the relata bound by a relevant law must be characterized partly in terms of hypothetical causal relations and partly in terms of actual causal relations to other things which are themselves characterized partly by actual and hypothetical causal relations. It is under that description that the syntax is nomically related to the target content. We don't know of any other laws like that.[2][3]

It is important to see that the case under consideration isn't like the standard functionalist case in which all the causal-functional constraints are met all at once by various concrete realizers, where all those laws target that single relatively concrete entity, and under that concrete characterization. In the case at hand, it is rather that some of the relations and roles create something as something to which the other causal relations can bind, but only when it is also bound by the first set. While (Burge 1986) and other externalist works have defended laws or at least explanations which link entities characterized relationally, they are not laws which link entities essentially characterized by hypothetical relations. And it is striking that Fodor himself has argued that even the non-hypothetically relational properties involved in broad content can't properly figure into psychological explanations. See (Fodor 1987, 27-53) and (Fodor 1991b).

Still, it may seem that even functionalism is at least implicitly committed to laws of this type. For instance, when something is functionally individuated as a headache, it is a headache because of its causal-functional role. But a headache may cause someone to go buy a pain reliever. So it looks like buying pain relievers enters into laws with headaches in virtue of the

functional roles brain states play in order to become headaches. But it is important to remember that a plausible form of AD must deploy nomic relations which link elite, projectible properties. And there is no proper law of this sort which links headaches and shopping for pain relievers.

Some forms of functionalism maintain that a higher-order, say psychological, functional structure may be realized by a variety of lower-order, say biological, functional structures. But that is not the sort of dependency within a level of structure which our modification of AD requires.

Perhaps it will seem helpful that there is a more concrete way to conceive the laws modified AD needs here. The laws might link target contents not to hypothetically and relationally characterized syntax, but rather to vast disjunctions across all the possible global internal physical states realizing such syntactic states, such that a content-fixing nomic relation is between the target content on one hand and a disjunction across all the concrete conditions including the relational ones sufficient for the syntax to be the syntax that it is on the other. But this suggests that the content doesn't accrue just to the intuitive syntax, but to the syntax plus the whole syntactic-typing physical environment of that syntax in the cognizer. And we will trace AD's necessary antipathy to disjunctive properties as properly elite and projectible in the next section. And in any case this is a very unusual and unfamiliar sort of nomic relation.

Whichever of these alternatives we adopt, disjunctive or staged causal-functional, no plausible law has this particular kind of complicated form. These stories might be made to work at the cost of enough questionable metaphysics. But we must keep our eye on the full measure of that cost. It is important to plausible AD accounts that the kind of laws which fix content involve genuinely elite and properly nomic properties, properties which among other things are projectible and support counterfactuals, and that many other kinds of intuitively causal relations are not elite in this way. The elite properties privilege some intuitive causal relations over others,

in the only plausible AD conceptions. But yet more is required by the proposal under consideration. There must be elite nomic relations among certain sorts of elite properties which together fix syntactic identity, and then a second set of elite nomic relations which link nomic structures of the first sort in a way to fix content, and indeed there probably must be dependencies among such second-level relations which help single out intuitive contents by the mechanism of AD. But if eliteness is projectibility, it is hard to see how there can be the two or three distinct sorts of elite properties this requires. At the very least, it must be explained how laws of the very peculiar kind AD would require on this option can be consistent with the truth of physicalism, how their truth can be in any sense constituted by, though perhaps not reduced, to the physical. For on the surface they are not. They seem to require platonic properties and platonic relations among those properties and then platonic relations of dependency among those relations which do not track anything like relatively concrete and familiar projectibility.

Dependency structures such as these deserve more attention than I can give them here. But they are unfamiliar. And it is also relevant that if, despite all appearances, they are available at acceptable metaphysical cost, they may provide a useful augmentation of internalist causal-functional accounts of content against the attacks of externalists like Fodor.

But in any case, we should remember that this has been consideration of an attempt which grants that modification of AD accounts is necessary. Our situation so far is this: If tokens are typed relatively concretely, that plausibly requires relatively narrow scope content-fixing laws, whose cost we will assess in the next section. If they are typed relatively abstractly, by reference to causal-functional relations, then AD fails, barring new forms of platonism or quasi-dualism or otherwise implausible claims about the nature of content-determining laws, which would in any case require some modification of AD.

In this section, we will be concerned primarily with problems for asymmetric dependence created by narrow scope typing. We will consider the interaction of narrow scope typing with the grounds issue, in particular by considering the other relata of the relevant content-fixing nomic relations, the target contents. The general problem we will discover arises out of the response required by one characteristic objection to AD accounts. This crucial response is problematic if content-fixing nomic dependencies are narrow in scope. And hence AD assigns unintuitive contents, barring some very implausible metaphysics.

Let's begin with the characteristic objection. It is at least a surface problem for AD accounts that they suggest counterintuitive contents in cases in which the nomic relation to the proper target property is mediated by a causal mechanism in such a way that that nomic relation seems asymmetrically dependent on the nomic relations linking the mental representation with intervening states of the mechanism. For instance, my mental representation of cows in my belief box may seem plausibly nomically related to cows only because it is nomically-linked with a range of retinal (or more generally proximal) stimulations which mediate my experience of cows. It would seem hence that AD accounts implausibly suggest that the content of such thoughts is a disjunction over retinal stimulations, and not cows.

It is worth considering Fodor's crucial response to this at length. No other response seems possible it AD is to be properly defended. Surely it wouldn't be adequate merely to <u>insist</u> on the contrary and AD-friendly dependency.

> It might still be said . . . that the dependence of cow thoughts on distal cows is asymmetrically dependent on their dependence on <u>disjunctions</u> of proximal cow projections; distal cows wouldn't evoke COW tokens but that they project proximal whiffs or glimpses or snaps or crackles or . . . well, or what? Since, after all, cow spotting can be mediated by theory to any extent you like, the barest whiff or glimpse of cow can do the job

> for an observer who is suitably attuned. Less, indeed, than a whiff or glimpse; a mere ripple in cow-infested waters may suffice to turn the trick. . . [J]ust about any proximal display might mediate the relation between cows and cow-thoughts for some cow-thinker or other on some or other cow-spotting occasion.
>
> So barring appeals to *open* disjunctions, it seems likely that there is just no way to specify an array of proximal stimulations upon which the dependence of cow-thoughts upon cows is asymmetrically dependent. And . . . it does seem to me that the price of intentional univocality is holding that primitive intentional states can't express open disjunctions. The idea might be that, on the one hand, content depends on nomic relations among properties and, on the other, nothing falls under a law by satisfying an open disjunction (open disjunctions aren't projectible) (Fodor 1990, 109-110). [24]

Such is Fodor's response, that the relevant disjunctions across proximal stimulations are open and anomic and unprojectible. But it is important to note that here the scope issue has become crucially relevant. This response requires that the content-determining laws be wide in scope. See for instance the last sentence of the first paragraph I've quoted. For each individual at each moment there *is* a particular mechanism present, specifying certain specific proximal stimulations which will generate mental representations of cows. Perhaps that is a disjunction across proximal stimuli, but hardly an open one. If relevant laws were very narrow scope, then it would only be because i) it is nomic for an individual at a time that "cow" is caused by the members of such a disjunction that ii) it is nomic for that individual at that time that cows cause "cow". So "cow" would mean the disjunction over proximal stimulations after all.

So if the laws invoked by AD are very narrow in scope, then this crucial response won't work. According to AD, "cow" would mean the disjunction across the proximal stimuli which would cause tokenings of "cow" in the individual in question at the time. AD would ascribe the

wrong content. But if, on the other hand, those laws are at all wide in scope, then we face the problem about syntax we discovered in the last section. According to AD, “cow” would mean typical syntactic cause of cow. AD would once again ascribe the wrong content. That is the crux of our dilemma for AD. It ascribes the wrong content whether the laws it invokes are wide or narrow in scope.

That is the main conclusion of this section, and indeed of the paper. But we should consider possible objections to this horn. It is here that our focus on the metaphysics of the grounds issue will eventually be relevant.

One might object that the difficulty at hand only attends a particular class of cases, so that while it is real it is of limited significance. AD might help in other cases. While it is of course impossible to consider all particular types of cases, and of course not all AD accounts are bound by Fodor’s own detailed commitments, still we can get some sense of the general force of my argument by considering all the cases which have been treated by Fodor in his role as AD’s primary partisan:

Perhaps the key and central cases of represented properties according to Fodor are what he calls “appearance” properties, which receive his latest extended attention in Concepts. Three examples are the properties of being red, being a doorknob, and being water (or at least being water for Homer). Fodor’s idea is that in the case of appearance properties, say the property of being a doorknob, things go like this: From experience of stereotypical doorknobs, we develop by a kind of generalization a mental representation of being a doorknob (and not just a stereotypical one). Indeed, he holds that it is essential to the property of being a doorknob that it is the property we “lock on to” from initial experience of stereotypical doorknobs, presumably by an innate mechanism and such that experience of stereotypical doorknobs can be noncircularly specified by reference to their concrete sensory properties. The three cases of appearance properties I’ve noted

are supposed to be somewhat distinct, in that the property of being a doorknob is not a natural kind, and in that specific sense doesn't cut nature at its joints. The property of being water is supposed to cut nature at its joints in that sense, though Fodor maintains that pre-chemical Homer didn't conceive of it as a natural kind, since (we can presume) he had no idea about inner essences and would have found the notion of water experts bizarre. Red is not a natural kind, but is unlike doorknob in that we can represent such a property in sensation even if we lack the concept, even if we can't represent it in thought. It is, Fodor says, a "sensory property". He explicitly holds, as indeed is plausible, that the property of being red and the property of being a doorknob are realized in the world itself by vast, and presumably largely unprojectible disjunctions over basic physical properties. He says indeed that the only law in which doorknobs qua doorknobs enter is the psychological law that we lock onto them in consequence of certain sorts of experience, namely of stereotypical doorknobs. (Fodor 1998a, 146).

Another important set of target contents are somewhat to the side of our main line of discussion, but also worth some consideration. Fodor plausibly maintains that the specifically logical vocabulary of Mentalese has its content fixed by a specification of its inferential role. (Fodor 1990, 110-111). And given predicate conjunction, he has argued following Evans, we can even help fix reference determinately in Quine's gavagai cases in this way, for instance distinguish between references to rabbits and undetached rabbit parts, and hence plausibly treat a set of cases which information-based semantics alone might otherwise be unable to secure. (Fodor 1994, 55-79). So to that degree inferential and hence causal-functional roles become relevant to what we might call the fine-grained referential structure of other mental representations than those which are logical words, along with the information-relevant nomic dependencies which have been our primary concern here.

So far Fodor's account seems largely consistent with internalist conceptions of the

grounds for content-fixing laws and dependencies. They are, as he says, laws about us. (Fodor 1998a, 146). But there are other cases. Fodor holds that we can come, by forming a conception of the inner essence of something like water and then deferring to experts about its nature, to have concepts of natural kinds as such. He says that, in such a case, by "water" we will, if in a largely H-O-H containing environment even with a bit of Putnam's XYZ thrown in and even if we haven't interacted with any H-O-H, refer to H-O-H and not XYZ. (Fodor 1990, 115-116). It seems to me that the most natural treatment of these cases for Fodor, modeled on his treatment of unicorns, is different from that he actual provides,[25] and makes much of the fact that we or at least our pre-chemical ancestors would have responded to XYZ as if to water if it had been around. But Fodor says that our deference to experts via AD yields the water-like substance our environment largely contains as the content of our mental representation of water, since in the largely H-O-H containing environment XYZ wouldn't cause tokenings of "water" unless H-O-H did. So this is externalism about the truth grounds for the relevant nomic and dependency claims.

Putnam's case of the elms and beeches suggests to Fodor that many of us lock on to the properties of being an elm and being a beech only through the mediation of our experts. (Fodor 1994, 33-39). He also holds that we lock on to the property of being a proton only via the mechanism of an experimental apparatus. (Fodor 1987, 118-126). In both cases there is a mechanism, containing experts or real machinery, which underwrites the particular nomic relations fixing the content. At one point Fodor also suggests that Mentalese names might refer by the use of an analogous mechanism, via historical community use as in Kripke, though this fits uneasily with his focus on synchronic dependencies and what he ways elsewhere about Davidson's swampman. (Fodor 1994, 115-119). In the case of the demonstratives of Mentalese, Fodor holds that their actual etiology is clearly relevant to their content. So here are more grounds for externalism.

That completes our survey of possible target contents which Fodor has treated. But notice that most of these cases generate a particular problem of the general sort under consideration, because they involve an intervening mechanism of a particular sort, at least if for argument's sake we accept the analysis of these cases which Fodor provides. The remaining cases, which don't involve such a mechanism, involve problems reminiscent of the difficulty with the relational typing of syntax we considered in section 2.

If laws have narrow scope, names in the usage of the individual in question would represent, according to AD and in light of the general problem presented by proximal mechanisms when laws are narrow in scope, the detailed and individual historical chains which mediate their reference and the usage of that individual. Demonstratives would present analogous difficulties. "Elm" and "beech" might represent one's particular experts saying "elm" and "beech". "Proton" might represent the particular blinking machinery which reveals or would reveal protons for someone. And the causing of the tokening of a mental representation by the property of being a doorknob or being red is plausibly mediated for a particular individual at a time by a specific set of proximal stimuli, and indeed can only be caused by a specific range of possible proximal stimuli. Indeed, there is a deeper layer of dependency here, on the stereotypical stimuli from which we would innately generalize to possession of the concept of doorknobs. If Fodor is right to say that being a property such that we lock to it on the basis of experience of stereotypical doorknobs is constitutive of what it is to be a doorknob, then he can hardly claim that it is obvious there is a mere diachronic dependency here. And he must admit that these various properties enter into psychological laws. So "doorknob" means the disjunction across the experiences of stereotypical doorknobs. If doorknobs are supposed to be lawlike because mind-dependent, surely it is worth noting that stereotypical doorknob looks also enter into psychological laws, and in fact into those very laws that are constitutive of what it is to be a

doorknob, according to Fodor.

We should also consider the case of the logical words, whose content is supposed to be fixed by conceptual and hence causal-functional roles. Presumably their syntax is fixed either by a weaker sort of causal-functional role or on a narrow scope basis by something like shape. Let me grant for argument that there are wide scope causal-functional characterizations of syntactic types and that we have suitably modified AD accounts to avoid the unintuitive alternative contents assigned by that relatively weak causal-functional role, in the manner suggested at the end of the last section. Still, note that AD assigns these logical words, by the very mechanism of their stronger and conceptually relevant causal-functional role, another sort of content in a way reminiscent of the problem raised in the last section. Their content is that which it is nomic would cause them and which is such that all other causes depend asymmetrically upon that, at least if we concessively rule any "first stage" syntax-fixing roles out of consideration. But then they are about the Mentalese words that cause them according to the conceptual-role elements of their strong causal-functional role. All that is required for this, even given the modifications to AD accounts introduced at the end of the last section, is that the relevant relations be nomic, and that other nomic dependencies, say mediated by the broad inferential role of such a phrase in a particular psychology, depend asymmetrically upon these.[26]

We can properly conclude that my argument, though initially focused on a single case, is of more than limited relevance. But let's consider alternative objections.

Fans of AD might try to wiggle up the middle on the scope issue. For instance, they might propose that content-fixing laws are invariant across language communities. But if the language of thought of an individual is constituted by their public language, then at least Fodor has lost perhaps his most characteristic thesis, and his arguments for it. More important, it is plausible that for a public language there are fixed distal stimuli which are the most characteristic causes

for appearance concepts and also the causes on which the other and more theoretical-mediated stimuli asymmetrically depend. So the content of “cow” in a public language, if that is fixed by AD, would probably be some disjunction across typical cows. And individuals speak quite differently, so it isn’t really clear that any token-typing of mental representations rooted in public language would be intrinsic anyway. So this middle thesis has some of the characteristic problems of both ends which surround it.

There is another salient middle ground thesis, that the relevant laws are laws for individuals over time. But remember that we cannot plausibly focus AD on diachronic dependencies among laws, and hence we might also properly wonder about the relevance of diachronic truth-grounds for laws. In fact, to speak about both diachronic and synchronic dependencies among laws seems to require that laws obtain at particular times, and that seems to imply that their truth grounds so exist. And in any case, it seems plausible that there is enough stability of belief in individuals to make it the case that there are laws of even this cross-temporal sort linking proximal stimuli disjunctions and “cow” for an individual. And it may be that very concrete syntactic typing is not available over time even for individuals.

So it seems that my argument cannot be evaded by tampering with the scope of the relevant laws. There are other sorts of response to the cow stimuli problem with which at least Fodor flirts, and hence which might seem to obviate the need for the alternative response which generates the problem I have suggested. For instance, Fodor says “I say that one might rule out proximal referents for mental representations by appeal to the principle that open disjunctions aren’t projectible. But one could also take the high ground and rule them out by stipulation: just as primitive symbols aren’t allowed to express necessarily uninstantiated properties, so too they aren’t allowed to express proximal properties.” (Fodor 1990, 134 n23). But of course this is not the high ground, but rather an obviously cheesy response, which is indeed always available to

defend any claim of sufficiency against any counterexample. And any specific exception like this means that, by stipulation, we can't represent such contents, at least directly and by ordinary means. What's more, this does involve some modification of AD accounts, and while perhaps it is plausible to defend an account of something which is very well intuitively-motivated by ruling out an occasional counterexample in this way, we saw in section 1 that AD accounts lack such motivation. And remember that the causal chain which links "cow" and cows for a particular individual includes lots of distal elements. So it might be impossible to create the relevantly complete list of things to be ruled out by stipulation as target contents, except by the unhelpful stipulation that everything which isn't the content doesn't count.

Another move which might be made here deploys another feature of Fodor's detailed accounts of AD, that mental representations of a certain type must not always be caused by their content, that they must sometimes be otherwise caused. (Fodor 1990, 118). What's wrong with deploying it here? The problem is that this additional constraint is met too easily for it to be effective, indeed by any plausible mental representation. Some "cow"s are not caused by the proximal array of sensory looks, but rather via inferences.[27]

There is another possible reply to the proximal stimulation problem for AD and hence to my argument which is suggested by Fodor's recent work on concepts. Fodor has recently argued that there are no recognitional concepts, concepts such that the ability to recognize at least some things that fall under the concept *as* things that fall under the concept is essential to possession of the concept. (Fodor 1998c and 1998d). This may seem the basis of a reply to my argument, because it apparently implies that a mental representation in an individual represents horses even when there is no particular set of stimuli and horsey looks which by law, even a law for that individual, would lead to it. The mental representation would be, for that individual at that time, nomically related to horses even were it not nomically related for that individual at that time to

horsey looks at all, even if there were no horsey looks which would lead the individual to recognize a horse.

But Fodor's argument is in fact unsuccessful. The argument against recognitional concepts rests on the claim that the ability to recognize good instances of a concept is not compositional, whereas concepts are. One can recognize good instances of pet fish without being able to recognize good instances of pets or of fish, and vice versa. This seems an odd argument for his conclusion, since the conclusion is that possession of a concept doesn't require an ability to recognize any instances, even in the case of perceptual concepts. (Fodor 1998c, 35). But, of course, if we can recognize a pet fish we can certainly recognize some pets and recognize some fish, even if not paradigmatic ones, and indeed if we can recognize some pets and some fish it is not implausible that we can recognize some pet fish, even if not paradigmatic ones, at least for all Fodor has argued. In addition, we can worry that Fodor's argument equivocates on "good instance", since whether something is a good instance of a concept depends on what kinds of things are generally in the environment in his central formulation of that notion (Fodor 1998d, 60), and yet at least in the case of sensory properties his argumentative target, the empiricists, will likely claim that relevantly good instances of red, the so-called unique reds, reds which are neither yellowish nor purplish, are fixed as such by individual neurophysiology or something else suitably internal and not by statistical analyses of instances of red in the environment.

There is one last response which AD might deploy to save at once narrow scope laws and intuitive contents, but it requires a very questionable position on the axis tracing possible grounds for claims of asymmetric dependence. Finally, we return to the grounds issue:

There might be an individual with a particular sort of physically-constituted causal structure which underwrites the truth of various counterfactuals and the falsehood of others, and also underwrites certain sorts of narrow-scoped projectibility claims rather than others, an

individual incorporating a mechanism which insures that in the face of a wildly disjunctive range of environmental stimuli such as those constituting the surface spectral reflecti normal humans see as red, mediated by a correspondingly disjunctive range of proximal stimuli, the individual's belief box will instance a particular mental word, say a token with a certain shape. What in the world makes it the case that it represents red and not the correspondingly disjunctive proximal stimuli? The physical facts about the individual don't seem sufficient to fix the story that AD needs to tell here if the laws have narrow scope. There are no individual and local physical resources to sort out the target properties as elite; the disjunction over proximal stimuli supports relevant counterfactuals and is locally projectible in the relevant way. And if the relevant laws have wide scope we run into the difficulty with typing which we faced in the last section.

Still, perhaps some metaphysics can help. It isn't enough to postulate objective platonic dependencies out in the world that constitute red as a nomic property from the point of view of the universe, if the laws are to be narrow in scope and the platonism is to be remotely consistent with physicalism. Rather, for there to be any help here from metaphysically elite properties, it seems that there would have to be a kind of special platonic idealization floating somewhere in the vicinity of each individual (or perhaps a somewhat wider-scoped target like a particular language community), involving special irreducible properties linked by irreducible nomic relations (ceteris paribus?) that fix that the proper idealization for semantic purposes of that particular individual (or community) focuses on cows rather than the proximal stimuli. This is close enough to dualism (epiphenomenal variety with supervenience and the causal-closedness of the physical) to give the game away, and might even be internalist dualism if the ghostly platonic mechanism is really a law about us. What AD would need to properly underwrite narrow scope laws which would deliver intuitive contents by asymmetric dependence only sounds at all plausibly consistent with a plausible physicalism if we ignore the issue of how physical truth is

supposed to ground all the very particular nomic and nomic dependency claims it needs to make to deliver intuitive contents.

The proper conclusion from the dilemma we have traced is that AD cannot play the central role in fixing content. This is not of course to conclude that it can play no role at all. Here are a few that still seem reasonable, for all I have argued: Assume that there are pieces of syntax determinately typed in an individual by something like shape. Then it might be that one causal-functional inferential role for that piece of syntax is asymmetrically dependent on another, so that asymmetric dependence helps to dispose of Fodor's primary objections to conceptual role semantics, by constituting one such inference as more deeply lawlike in that psychology than another. Alternatively, perhaps there is some range of sensory stimuli, even of phenomenal states, which are more deeply the causes of certain mental representations in the belief-box of an individual than certain others, such that the others asymmetrically depend on that range. Fodor seems (I think properly) unimpressed that recognitional accounts of concepts are deployed in anti-cartesian skeptical arguments. (Fodor 1998c, 35-36). So perhaps in this spirit we could presume that sensory states have some sort of content fixed independently of their environment. The resources introduced by AD accounts may help deliver an internalist, narrow scope account of content governed by a robust form of methodological solipsism, though they cannot play the central role in answering error or disjunction problems to externalist information-based accounts of content. Ironic, isn't it? Jerry Fodor, come home.[28]

Department of Philosophy
University of Nebraska--Lincoln

Notes

1. Canonical statements are perhaps (Fodor 1987, 97-127) and (Fodor 1990).

2. AD accounts have received significant critical discussion, in (Cummins 1989), (Maloney 1990), (Antony and Levine 1991), (Baker 1991), (Boghossian 1991), (Adams and Aizawa 1992, 1993, 1994, and 1997), (Manfredi and Summerfield 1992), (Seager 1993), and (Wallis 1995). But much of this discussion has focused on forms of AD which deploy mere counterfactual conditionals rather than the more robust nomic relations which Fodor has preferred in more recent formulations. I will sketch new problems which cannot be evaded by this reformulation.

3. For instance, some information-based accounts deploy nomic relations which assure that Y *would* cause X under certain hypothetical conditions, while others focus on *actual* and relatively accidental causal histories, as we will shortly see. For other related refinements see (Adams and Aizawa 1992, 1993, and 1994).

4. For discussion of the varieties, see (Manfredi and Summerfield 1992).

5. Perhaps (Fodor 1990, 93) is canonical.

6. In Fodor's terminology, nomic relations differ from actual causation in a second important respect to which we will return.

7. For instance, see (Cummins 1989) and also (Adams and Aizawa 1994, 239-240).

8. Fodor has occupied more than one of these argumentative positions over time.

9. See note 2 for references.

10. One crucial text is (Fodor 1990, 101-102). See also (Fodor 1990, 109) and (Fodor 1991a, 257).

11. It is also relevant in this context that, as indicated in note 2, much of the criticism which AD accounts initially faced focused on conceptions of AD which deploy counterfactual dependencies among causal relations, and not the more metaphysically robust conception of the nomic which Fodor endorses. See for instance (Cummins 1989), (Maloney 1990), (Boghossian 1991), (Manfredi and Summerfield 1992). The standard semantics for counterfactual conditionals deployed in this critical work may indeed force in a second way that Fodor deploy his robust metaphysics of nomic relations.

12. Another possible axis concerns the role in which an AD account is deployed, as an actual account of mental content or as a merely possible account. Fodor's response to certain objections that the nomic dependencies he claims are implausible, and in fact that even some of his own concrete cases suggest other treatments which would according to an AD account involve an implausible assignment of contents, is that he is doing philosophy and trying to show that content *can* be naturalized and hence he gets to specify the relevant counterfactuals and dependencies. See (Fodor 1990, 96). He isn't looking for *actual* grounds for content but only in some sense *possible* grounds, for sufficient but not necessary or even plausible conditions. But I will presume that we are interested in what actually constitutes thought content.

13. In fairness, there is a prevailing Fodorian view on this topic, underwritten by Fodor's general endorsement of apparently wide-scope psychological laws, though *ceteris paribus*. But of course AD accounts in general are not bound by Fodor's individual commitments, and we will see in section 2 that there is more than a little reason for an AD account to favor narrow scope laws.

14. Fodor has suggested several particular accounts of mental syntax, and there is even a dominant suggestion in his writings which is perhaps necessitated by his customary commitment

to wide scope laws. But the particular accounts he suggests are different at different times, and we will need to see how each possibility interacts with the other issues which concern us.

15. Of course, it might be supposed that there are two sorts of psychological laws relevant to content, some narrow-scope and some wide-scope, which generate two corresponding sorts of content, narrow and wide. But each alternative would suffer from one of the two problems we will trace.

16. This seems to be the primary drift of Fodor's thought on this matter. See the relatively recent statement in (Fodor 1994, 108-109), though we might worry about the use of numerical identity in that passage. In light especially of his most recent discussions of syntax, in (Fodor 2000, 20-22 and 28-33), these causal-functional roles would need to be relatively narrow, in other words they would need to be roles relevant to logical form but not to conceptual role in any very robust sense. While on page 20 Fodor suggests that syntactic properties are "local", this means only that the syntactic identity of a complex is determined by the syntactic identity and arrangement of its syntactic constituents. So this isn't reason to think his dominant account has shifted towards something concrete like shape.

17. On balance, I think this is Fodor's position.

18. Fodor, given his support of punctate minds, might implausibly disagree. But in any case he faces the other problems noted.

19. This is somehow shared in all of its occurrences, be they in belief or desire boxes. There are allied views on which the belief box tokens are typed in one way and there is an alternative form of typing which links different tokens playing different roles within a given individual.

20. In fact, Fodor's own discussion of causal-functional typing in the relatively recent (Fodor 1994) adverts to a machine table which seems to be an *individual* machine table, or at least a table for a number of machines more similar than any two humans are.

21. Fodor himself sometimes contemplates another move against counterexamples which may seem helpful at this point: Bald stipulation. See (Fodor 1990, 134). He might deploy a bald stipulation that the relata which are the content can't be among those relevant to syntactic typing. We will return to this desperate response in the next section.

22. See (Kim 1993) and Fodor's response in (Fodor 1998b).

23. Perhaps some hold that all basic laws relate pure dispositional properties, but then such an account is by that very reason problematic.

24. See also (Fodor 1991a, 257).

25. In fact, the treatment of twin cases in (Fodor 1994) claims that they are nomologically impossible and hence irrelevant.

26. Notice that this problem involving conceptual roles will infect even an AD account that deploys syntactic typing by shape.

27. This discussion is related to (Anthony and Levine 1991). Fodor's response to their discussion is that the dependence of the mental representation on the disjunction is diachronic, and that the proper asymmetric dependencies are synchronic. See (Fodor 1991a, 313 n1). But this isn't adequate. The question in an AD account is not about dependencies among actual causes but among nomic relations, and it is implausible that a mental representation would still mean horse

if horsey stimulations wouldn't cause it in the moment. And remember the synchrony of the relevant dependencies.

28. Thanks for comments from John Gibbons, Leo Iacono, J.D. Trout, Mark van Roojen, and two anonymous referees.

References

Adams, Frederick and Kenneth Aizawa. (1992) "'X' means X: semantics Fodor-style", Minds and Machines 2, 175-183.

Adams, Frederick and Kenneth Aizawa. (1993) "Fodorian semantics, pathologies, and Block's problem", Minds and Machines 3, 97-104.

Adams, Frederick and Kenneth Aizawa. (1994) "Fodorian Semantics", in Stephen P. Stich and Ted A. Warfield (eds.), Mental Representation, Cambridge, MA, Blackwell, 223-242.

Adams, Frederick and Kenneth Aizawa. (1997) "Fodor's Asymmetric Causal Dependency Theory and Proximal Projections", The Southern Journal of Philosophy 35, 433-437.

Antony, Louise and Joseph Levine. (1991) "The Nomic and the Robust", in Barry Loewer and Georges Rey (eds.), Meaning in Mind: Fodor and his Critics, Cambridge, MA, Blackwell, 1-16.

Baker, Lynne Rudder. (1991) "Has Content Been Naturalized?", in Barry Loewer and Georges Rey (eds.), Meaning in Mind: Fodor and his Critics, Cambridge, MA, Blackwell, 17-33.

Boghossian, Paul. (1991) "Naturalizing Content", in Barry Loewer and Georges Rey (eds.), Meaning in Mind: Fodor and his Critics, Cambridge, MA, Blackwell, 65-86.

Burge, Tyler. (1986) "Individualism and Psychology", Philosophical Review 95, 3-45.

Cummins, R. (1989) Meaning and Mental Representation, Cambridge, MA, The MIT Press.

Fodor, Jerry. (1980) "Methodological solipsism considered as a research strategy in cognitive science", Behavioral and Brain Sciences 3, 63-109.

Fodor, Jerry. (1987) Psychosemantics, Cambridge, MA, The MIT Press.

Fodor, Jerry. (1990) “A Theory of Content, II: The Theory”, A Theory of Content and Other Essays, Cambridge, MA, The MIT Press, 89-136.

Fodor, Jerry. (1991a) “Replies”, in Barry Loewer and Georges Rey (eds.), Meaning in Mind: Fodor and his Critics, Cambridge, MA, Blackwell, 255-319.

Fodor, Jerry. (1991b) “A modal argument for narrow content”, Journal of Philosophy 88, 5-26.

Fodor, Jerry. (1994) The Elm and the Expert, Cambridge, MA, The MIT Press.

Fodor, Jerry. (1998a) Concepts, Oxford, Clarendon Press.

Fodor, Jerry. (1998b) “Special Sciences: Still Autonomous after All These Years”, In Critical Condition, Cambridge, MA, The MIT Press, 9-24.

Fodor, Jerry. (1998c) “There are No Recognitional Concepts--Not Even RED”, In Critical Condition, Cambridge, MA, The MIT Press, 35-47.

Fodor, Jerry. (1998d) “There are No Recognitional Concepts--Not Even RED, Part 2: The Plot Thickens”, In Critical Condition, Cambridge, MA, The MIT Press, 49-62.

Fodor, Jerry. (2000) The Mind Doesn’t Work That Way, Cambridge, MA, The MIT Press.

Kim, Jaegwon. (1993) “Multiple Realization and the Metaphysics of Reduction”, Supervenience and Mind, Cambridge, Cambridge University Press, 309-335.

Maloney, J. Christopher. (1990) “Mental Misrepresentation”, Philosophy of Science 57, 445-458.

Manfredi, Pat A. and Donna M. Summerfield. (1992) “Robustness without Asymmetry: A Flaw in Fodor’s Theory of Content”, Philosophical Studies 66, 261-283.

Seager, William. (1993) “Fodor’s Theory of Content: Problems and Objections”, Philosophy of Science 60, 267-277.

Wallis, Charles. (1995) “Asymmetric Dependence, Representation, and Cognitive Science”, The Southern Journal of Philosophy 33, 373-401.