



Special Sciences: Still Autonomous After All These Years

Author(s): Jerry Fodor

Reviewed work(s):

Source: *Noûs*, Vol. 31, Supplement: Philosophical Perspectives, 11, Mind, Causation, and World (1997), pp. 149-163

Published by: [Blackwell Publishing](#)

Stable URL: <http://www.jstor.org/stable/2216128>

Accessed: 12/11/2011 16:57

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at

<http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Blackwell Publishing is collaborating with JSTOR to digitize, preserve and extend access to *Noûs*.

<http://www.jstor.org>

**SPECIAL SCIENCES:
STILL AUTONOMOUS AFTER ALL THESE YEARS***

Jerry Fodor
Rutgers University

The conventional wisdom in philosophy of mind is that “the conventional wisdom in philosophy of mind [is] that psychological states are ‘multiply realized’ ...[and that this] fact refutes psychophysical reductionism once and for all. (Kim, 1992; p. 1. All Kim references are to this paper except as noted.)” Despite the consensus, however, I am strongly inclined to think that psychological states are multiply realized and that this fact refutes psychophysical reductionism once and for all. As e.e. cummings says somewhere: ‘*nobody* loses *all* of the time.’

Simply to have a convenient way of talking, I will say that a law or theory that figures in bona fide empirical explanations, but that is not reducible to a law or theory of physics, is ipso facto *autonomous*; and that the states whose behavior such laws or theories specify are *functional* states. (In fact, I don’t know whether autonomous states are ipso facto functional. For present purposes all that matters is whether functional states are ipso facto autonomous.) So, then, the conventional wisdom in the philosophy of mind is that psychological states are functional and the laws and theories that figure in psychological explanations are autonomous.¹ (Likewise, and for much the same reasons, for the laws, theories, etc. in other ‘special’ (viz, nonbasic) sciences.) The present discussion undertakes to defend this conventional wisdom against a philosophical qualm that I take be the main moral of Jaegwon Kim’s recent paper “Multiple realization and the metaphysics of reduction.”

Kim says that he’s prepared to agree (at least for the sake of the argument) that:

1. Psychological states are typically multiply realized (MR);

and

2. MR states are ipso facto unsuitable for reduction.

But Kim thinks philosophers haven't gotten it right about *why* MR states are ipso facto unsuitable for reduction. Once they do, Kim says, they'll see that the moral of 1&2 isn't, after all, that psychology is autonomous. Rather, it's that quotidian psychological states aren't *reducible* because they aren't *projectible*. Unprojectible states are, by definition, not the subjects of a possible science; they aren't bona fide kinds and they can't appear in bona fide nomological explanations. A fortiori, terms that express psychological states are not available for incorporation in "bridge laws" or in (metaphysically necessary) property identities. This is all, of course, contrary to what a lot of philosophers, to say nothing of a lot of psychologists, have hitherto supposed.

Caveat: It turns out, according to Kim, that some very finely individuated psychological states are 'locally' reducible; but that's because these states are, by assumption, not MR. More on this later; I mention it here just to make clear that Kim is *not* claiming that psychological states are unprojectible *qua psychological* (or *qua intentional*), but only that they are unprojectible *qua not local*. Contrast (eg.) Davidson, (1980).

Now, I think Kim is quite right that what moral you should draw from MR states not being reducible depends on whether whatever it is that makes them not reducible also makes them not projectible. But I think that the diagnosis of the irreducibility of MR states that Kim offers is wrong, and that the right diagnosis supports the consensus conclusion: 'pain', 'believes that P' and the like express real states, about which all the available evidence suggests that there are real laws.² If I'm right that such states are projectible but not reducible, then it follows that psychological laws are autonomous after all.

So much for strategy, now for tactics. Kim's polemical method is to pick an apparently untentious (and, in particular, a *nonpsychological*) MR state and explain why it is unprojectible *qua* MR. He will then argue that, since beliefs, pains and the like are MR by assumption (1), they must be unprojectible for the same reasons that his MR paradigm is. Tit for tat: I'll argue both that Kim's diagnosis of the unprojectibility of his MR paradigm is wrong and that the supposed analogy between his MR paradigm and pains, beliefs and the like is spurious.

More caveat: In order not to be always writing 'pains, beliefs and the like' or 'quotidian psychological states,' I'll follow Kim's practice and take pain as my working example of a quotidian psychological state: In particular, I'll assume that, if there are any psychological laws at all, then probably there are psychological laws about pain. In fact, however, pain isn't really a happy choice for a working example since, though I take it to be quite a plausible candidate for projectibility, pain is notoriously a very bad candidate for being MR, hence for functional analysis. That, however, is stuff for a different paper. (See, eg. Block and Fodor, 1972; Block, 1978). My present brief is just that what Kim says is wrong with functionalism isn't; so the autonomy thesis is, to that extent at least, not in jeopardy. For these purposes, pain will do.

So, finally, to work.

Jade

We are told, Kim tells us, that jade “is not a mineral kind, contrary to what was once believed; rather, jade is comprised of two distinct minerals with dissimilar molecular structure, *jadeite* and *nephrite*. (11)” Geological unsophisticate that I am, I shall often call these two minerals jadeA and jadeB in the discussion that follows. It won’t really matter which is which, so you needn’t bother to keep track.

Kim thinks that, because of these facts about jadeite and nephrite, jade is paradigmatically MR. Kim also thinks that since “jade” is paradigmatically MR, it is ipso facto unprojectible. And, a fortiori, (3) isn’t a law.

3. Jade is green.

I don’t actually care much whether (3) is a law or even whether “jade” is projectible. But I am going to deny that jade is paradigmatically MR, and I am also going to deny that jade is unprojectible for the reason that Kim says that it is.

So then, what, according to Kim, is wrong with (3)? Well, “(l)awlike generalizations...are thought to have the...property [that] observation of positive instances, Fs that are Gs, can strengthen our credence in the next F’s being G.” In short, real laws are confirmed by their instances, but (3), according to Kim, is not.

“...we can imagine this: on re-examining the records of past observations, we find, to our dismay, that all the positive instances of (3)...turn out to have been samples of jadeite, and none of nephrite! If this should happen...we would not...continue to think of (3) as well confirmed...But all the millions of green jadeite samples *are* positive instances of (3)...the reason [that (3) is not confirmed by them] is that jade is a true disjunctive kind, a disjunction of two heterogeneous nomic kinds which, however, is not itself a nomic kind.(12)”

Now notice, to begin with, that the thought experiment Kim proposes doesn’t really make the case that he wants it to. To be sure, if we discovered that all our samples of green jade are samples of green jadeA, that would lead us to doubt that (3) is well confirmed. Not, however, because it would be a reason to think that (3) doesn’t “pass the projectibility test (12)”; rather because it would show that we’ve made a sampling error in collecting the data. The point is: *Anybody* can make a sampling error, *whether or not* the hypothesis he’s trying to confirm is projectible.

Suppose we’ve been considering whether oak trees shed their leaves in winter; and suppose it turns out, on re-examining the records, that all our positive instances are observations of oak trees on the north side of hills. Then we would no longer think of the generalization about oak trees losing their leaves in the winter as unambiguously well-confirmed; oak data confirm oak generalizations

only if they are an *unbiased* sample of the oak population; which, on the current assumption, our data aren't. Maybe, in the present case, the generalization that the instances really confirm is that *oak trees on the north side of hills* lose their leaves in winter.

But notice that discovering a sampling error of this sort would be no reason at all for doubting that 'oak tree' is a kind. Rather, the worry would be that maybe *oak tree on the north side of a hill is a kind too*. If it is, then our data are equivocal between two perfectly ok, projectible hypotheses: the one that goes *blah, blah...oak trees...* and the one that goes *blah, blah.... oak trees on the north sides of hills...* . When we discover the sampling error, we regard neither generalization as unequivocally confirmed by data that are instance of both, and this is precisely because the data *are* instances of both. *The sampling error means that the data are equivocal, not that the hypotheses are unprojectible*. There is, to be sure, something wrong with (3); something that makes it not a law. But what's wrong with (3) isn't that biased samples fail to confirm it. *Biased* samples don't confirm *anything*.

So Kim's thought experiment shows nothing that is to the point. However, he doesn't really need the thought experiment to run his argument. Instead, he could run the argument like this:

Although green JadeA samples tell us that jadeA is green, and green jadeB samples tell us that jadeB is green, green jadeA samples tell us nothing at all about the color of jadeB and green jadeB samples tell us nothing at all about the color of jadeA. That's because, though it's true that jadeA is green iff jadeB is green, that it is true is merely accidental given the presumed facts about the structural heterogeneity of nephrite and jadeite. Analogously: Since jadite is green and grass is green, jadeite is green iff grass is. But it doesn't follow that evidence about the color of grass bears at all on hypotheses about the color of jadite or vice versa.

But, now, according to the functionalist orthodoxy, pain is just like jade, isn't it? So, then: "Why isn't pain's relationship to its realization bases, N_h , N_r , N_m analogous to jade's relation to jadeite and nephrite? If jade turns out to be nonnomic [i.e. not projectible] on account of its dual 'realizations' in distinct microstructures, why doesn't the same fate befall pain? After all, the group of actual and nomologically possible realizations of pain, as they are described by the MR enthusiasts with such imagination, is far more motley than the two chemical kinds comprising jade... . we put the following question to Fodor and like-minded philosophers: If pain is nomically equivalent to N, the property claimed to be wildly disjunctive and obviously nonnomic, why isn't pain itself equally heterogeneous and nonnomic as a kind? (15; quoted emphasis Kim's)".

I expect this question is meant to be rhetorical, but I think that I'll answer it all the same.

Rebuttal

There is, I think, a sort of polemical standoff at this point in the discussion: Kim is quite right that jade is “a true disjunctive kind, a disjunction of two heterogeneous nomic kinds which, however, is not itself a nomic kind.(12)” But, from the functionalist’s perspective, to say that jade is a disjunctive kind because jade generalizations aren’t confirmed by their instances is to put the epistemological cart before the ontological horse. The right story, according to the functionalist, goes the other way around. What makes jade a disjunctive kind (in the, as we’re about to see, special sense that is germane to whether it’s projectible) is that *there are no general empirical truths about jade as such*; a fortiori, there are no such truths for samples of jade, as such, to confirm. Whatever is reliably true of jade is reliably true of jade either because it is reliably true of *jadeA* as such or because it is reliably true of *jadeB* as such.

No doubt there are those of you who are suspicious of ‘as such’s as such; but, in fact, none are required to make the present point. What’s needed is just the distinction between a multiply based property that is *disjunctive*, and a multiply based property that is disjunctively *realized*. To wit: A multiply based property is disjunctive iff it has no realizer in *any* metaphysically possible world that it lacks in the *actual* world. Jade is disjunctive because the only metaphysically possible worlds for jade are the ones which contain either jadeA, or jadeB or both. By contrast, multiply based properties that are disjunctively *realized* have different bases in different worlds. Pain is disjunctively realized because there’s a metaphysically possible, nonactual, world in which there are silicon based pains.

That is to say, in effect, that being jade *just is* being jadeA or jade B. Whereas it would be simply question begging of Kim to hold that being pain is the property of being one or other of pain’s realizers. Functionalists claim that pains and the like are higher-order, relational properties that things have in virtue of the pattern of causal interactions that they (can or do) enter into. *If so, then pains, though multiply based, are not disjunctive but MR.*

To repeat: Though Kim *says* that he concedes that psychological properties are MR, that’s only because he isn’t distinguishing *being MR* (like pain) from *being disjunctive* (like jade). But it’s exactly the distinction between *disjunctiveness* and disjunctive *realization* that functionalists are insisting on when they say that pain states are nomologically homogeneous under their functional description *despite the physical heterogeneity of their realizers*. You can’t (Kim can’t) refute this claim just by defining ‘disjunctive kind’ so that it isn’t true.

This is, as I say, a sort of polemical standoff. The functionalist assumes that there are laws about pains ‘as such’; so he infers that, though pain is multiply based, it is *not* (merely) disjunctive. So he infers that pain is unlike jade in the respects that are relevant to the question of projectibility. Kim, going the other way around, assumes that pain is (merely) disjunctive, hence that it is relevantly similar to jade, and hence that there aren’t any laws about pain. The real issue—

the one that Kim's appeal to the jade example begs—is whether there *is* a difference between a multiply based property's being disjunctive and its being MR; and, if so, what the difference is and whether it matters to projectibility.

Kim almost sees this in the closing sections his paper. But then he gets it wrong; fatally, in my view. “Is jade a *kind*? We know it is not a mineral kind; but is it any kind of a kind... There are certain shared criteria, largely based on observable macroproperties of mineral samples (e.g. hardness, color, etc.), that determine whether something is a sample of jade... . What all samples of jade have in common is just these observable macrophysical properties that define the applicability of the predicate ‘is jade’. (24)” This, I say, is just wrong; and resolving the metaphysical issues about projectibility that Kim has raised turns on seeing that it is.

Suppose that, pottering around in the cellar one day, you succeed in cooking up a substance—out of melted bottle glass, let's say—that is, for all macroscopic purposes, indistinguishable from jade: For example, it's as similar in color to jadeA and jadeB as they typically are to one another; its hardness falls at about the right place between talc and diamond on the scratch test; it cracks along the right sort of cleavage planes; it weighs about the same as jade per unit volume ... and so forth. Have you, then, *created jade*? Oh frabjous day! Oh joy that alchemists never knew! Oh (in particular) riches running wild!

Not on your Nelly. What you've got there isn't jade; it's just melted bottle glass. Melted bottle glass maybe counts as *artificial* jade in the sort of case that we've imagined; but *do not try to sell it as the real stuff*. You will find, if you do, that fanciers of jade are not amused. They will call you unkind things like *thief* and *fraud*; and if they catch you, they will lock you up.

Pace Kim, being jade is *not* relevantly like having a functional (i.e. MR) property; if it were, you could make new jade by mimicking the macroscopic properties that jadeA and jadeB share (the ones which, according to Kim, “determine whether something is a sample of jade, or whether the predicate ‘is jade’ is correctly applicable to it.”) But you can't make jade that way. If you want to make some jade, you have to make either some jadite or some nephrite; just as, if you want to make some water, you have to make some H₂O. That's because *jade is jadite or nephrite is metaphysically necessary*, just like *water is H₂O*.³ As with most of the metaphysical claims one comes across these days, the one that I just made relies for its warrant on a blatant appeal to modal intuitions. But I think the modal intuitions that I'm mongering are pretty clearly the right ones to have. If you don't share mine, perhaps you need to have yours looked at.

Now compare pain according to the functionalist persuasion. Suppose you should happen, one day down in the cellar, to throw together a robot; among whose types of internal states there's one that is, under functional description, about as similar to my pains as my pains are to yours. Then, functionalists say, the thing that you've created *has pains*. Not *artificial* pains, not pain *simulations*, not *virtual* pains, but *the real things*. Even if, as we may suppose, what you've cooked up is a silicon robot that's made out of melted bottle glass.

My point is not, I emphasize, to claim that the functionalist is right to be of this persuasion. For reasons I've elaborated elsewhere, I doubt that pain is a functional kind, or that beliefs-that-P are either. What's at issue, remember, isn't whether functionalism is true; it's only whether there is a plausible functionalist response to the challenge that Kim has laid down: "Why isn't pain's relationship to its realization bases...analogous to jade's relationship to jadeite and nephrite?" Reply: There's a difference between being a functional property (being multiply realized) and being a disjunctive property. Being jade, according to the geologists, is an example of the latter; being pain, according to the psychologists, is an example of the former. So there is, thus far, nothing to warrant an inference from the unprojectibility of jade to the unprojectibility of pain. Kim is quite right to emphasize that "the group of actual and nomologically possible [and epistemically and metaphysically possible] realizations of pain, as they are described by the MR enthusiasts with such imagination is far more motley than the two chemical kinds compromising jade." What he's missed is that *for that very reason* the pain/jade analogy is flawed and won't do the polemical work that he wants it to.

If there is after all plausibly a difference between the relation between pain and its realizers, on the one hand, and the relation between jade and its realizers, on the other, then it's not patently irrational of a functionalist to agree that 'jade' isn't projectible while continuing to claim that 'pain' is. But what is not patently irrational may nonetheless be unreasonable and obstinate. Granted that pain and jade differ in the ways that we have been exploring, it remains open to Kim to wonder why that kind of difference should matter so much. How come is it that MR properties are projectible are disjunctive properties aren't?

Or, to put the same question slightly differently: Functionalists are required to deny that pain is *identical* to the disjunction of its realizers. The reason they are is that it's part of their story that the functional property realized, *but not its physical realizer*, is projectible. And the reason they have to say *that* is that *otherwise multiple realization wouldn't be an argument against reduction*: What is supposed to make the case for the autonomy (/unreducibility) of functional laws is that there aren't any laws about the realizer of a functional state *even if there do happen to be laws about the functional state that they realize*. But then functionalists must themselves think that disjunctions—including, notice, the disjunctions that realize bona fide MR states—aren't projectible. What justifies them in making this claim?

Or, to put the question yet another way: Functionalists hold that the biconditionals that connect functional properties with their realizers aren't *laws* (a fortiori, that they aren't 'bridge' laws.) They *can't* be laws because the realizers of functional states are, by assumption, disjunctive; and disjunctive properties are supposed to be *ipso facto* not projectible. Thus, in one of his earliest functionalist papers, Putnam remarks that rendering realizers projectible by "...defining the disjunction of two [realizing] states to be a single 'physical-chemical state'..." is not a metaphysical option that can be taken seriously." (Rosenthal (1991, p. 201)). I think Putnam is right that it's not. But the question *why* it isn't remains.

Here's where we've got so far: Functionalists think that there are no metaphysically open laws (A law about a multiply based property is 'metaphysically open' iff that property is MR rather than (merely) disjunctive; i.e. iff there are metaphysically possible worlds in which it has realizers that it doesn't have here). Accordingly, functionalists think that there are laws about pain, but they don't think that there are laws about jade, *and they also don't think that there are laws about the metaphysically open disjunctive realizer of pain*. So functionalists agree with Kim that disjunctions aren't projectible, *whether they are metaphysically open or metaphysically closed*. But, pace Kim, the reason open disjunctions aren't projectible is *not* that their instances fail to confirm them. So what is the reason? In one form or another, this question keeps coming up. It needs an answer.

Here's a possible story: Suppose, for a moment, that metaphysically open disjunctive properties are projectible after all. Still, there's a difference between the case where a disjunction appears *only* in bridge biconditionals, and the case where it is, as it were, independently certified because it also occurs in laws at its own level. It might be quite reasonable to hold that disjunctions are bona fide in bridge laws *only* when their projectibility is independently certified in this way. That would distinguish 'real' bridge laws from those containing formulas that are not only disjunctive, but also *gerrymandered*; the latter being cases where, intuitively, *all* that the disjuncts have in common is that they realize some higher level state.

That was, in effect, the view I took in an earlier paper (also called "Special sciences" (Fodor, 1974)). The objection I voiced there was not to the projection of disjunctions—open or closed—as such; but rather to a theory of interlevel nomic relations that fails to distinguish real reductions from gerrymanderings. 'Type' physicalism is supposed, by general consensus, to be stronger than 'token' physicalism; stronger, that is, than the mere claim that all mental states are necessarily physically instantiated, however unhomogeneously. I suggested that the difference is that type physicalism, but not token physicalism, requires bridge laws that really are laws; viz ones that contain predicates that are independently certified as projectible because they are independently required to state *intralevel* laws. By this 'no gerrymandering' criterion, it's empirically plausible that the multiple realizers of functional states are often not projectible, either in psychology or elsewhere.

I still think that's perfectly alright, as far as it goes. In fact, I think it's quite attractive since it fends off the following dilemma. I suppose a functionalist might wish to admit that there are nomologically necessary constraints on what sorts of things can be realizers of pains (or canopeners or whatever). That is, for each functional property P, there must be *some* disjunction of realizers (call it '#') such that 'Ps are R1, or R2, or R3...#' is nomologically necessary;⁴ and even if we don't (/can't) know what # is, presumably God can (/does). But if 'Ps are R1, or R2, or R3...#' is nomologically necessary it looks like P is reducible after all since it looks like 'Ps are R1, or R2, or R3...#' is a (bridge) law; presumably, a law is

just a universal conditional that's true in all nomologically possible worlds. But the argument for the autonomy of metaphysically open MR states depended on their *not* being lawfully related to their realizers. So it looks like that argument must be unsound.

The way out of this for functionalists is to require that bridge laws be not just nomologically necessary but also that they be *not gerrymandered*. What's wrong with 'Ps are R1, or R2, ...' isn't that it lacks a nomologically necessary closure; '#' is one by assumption. What's wrong is rather that *the nomological closure is gerrymandered*. I.e. 'R1, or R2 or... , #' isn't 'independently certified'. I.e. it doesn't occur in any *proper* ('single level') laws. Since 'R1, or R2, or... , #' isn't independently certified, 'Ps are R1, or R2, or... , #' isn't a bridge law and (*ceteris paribus*) P isn't reducible after all. On this account, the constraints on bridge laws are *stronger than* (in effect, they include) the constraints on proper (single-level) laws. This difference is what underlies the intuition that type materialism comes to more than just the claim that it's nomologically necessary that every nonbasic property be physically realized.

As I say, I still like that story; but I admit that there may be more to be said. It's clearly Kim's intuition that there's something wrong with multiply based kinds as candidates for projectibility *whether or not* they are gerrymandered. I.e. that you aren't allowed multiply based kinds in either inter- or in intra-level laws. For reasons we've already discussed, I don't accept Kim's account of why this prohibition holds; on my view, his diagnosis depends on his failure to distinguish closed disjunctions from multiple realizations. But I'm prepared to take it seriously that maybe disjunctions are, as such, bad candidates for projection. And as we've just seen, even functionalists have to claim, at a minimum, that nondisjunctive functional states are ipso facto better candidates for projection than are their open disjunctive realizers; otherwise open disjunctive laws about the latter would compete with nondisjunctive laws about the former. So, then, *just what is wrong with projecting disjunctions?*

It's not hard to see why it's so plausible that there can't be laws about *closed* disjunctions. Presumably the nomic properties that a thing has qua F or G are either properties that it has qua F or properties that it has qua G. That's why, if being jade really is a disjunctive property (if being jade is just being jadite or nephrite) then *of course* there are no laws about being jade 'as such'; all the jade laws are ipso facto either jadite laws or nephrite laws.

But it's not so obvious why (/whether) there can't be laws about *open* disjunctions as such. In fact, I think there are depths here to be plumbed.

One might, for starters, try denying that open *formulas*—ones that contain ellipses essentially—succeed in expressing properties at all; in which case, their failure to express projectible properties would not be surprising. But I propose to assume without argument that a predicate like 'is $R1 \vee R2...$ ' does pick out a corresponding property; viz the property of *being* $R1 \vee R2...$. It is, remember, common ground between Kim and the functionalists that mental states are mul-

tively realized. So both are committed to some sentence of the form 'pain is $R1 \vee R2...$ ' being true; and it's hard to see how one could be unless predicates with ellipses can refer.

To take for granted that (openly) disjunctive sentences (can) have truth conditions is not, of course, to say what their truth conditions are. I haven't got a semantics for such sentences to offer you; and trying to construct one is not a task that I relish. Maybe they should be treated as true at each world where any of the disjuncts is true, and as neither true nor false (as noncommittal, if you like) anywhere else. In what follows, I'll assume some story of that sort; I doubt the details matter much for the polemical purposes I have in mind.

Where we are now is: We're assuming that open disjunctions can express bona fide properties, yet that the properties that they express are somehow *in-trinsically* unfit for projection. This can really seem quite puzzling. Presumably God can do anything He likes with the properties He has at hand; so if there really are such properties as *being* $R1 \vee R2...$, why can't God make laws about them?

Nonetheless, I think the intuition that open disjunctions are at best bad candidates for laws is basically sound. Here's why: *Open laws suggest missed generalizations*. To offer a law of the form $R1 \vee R2 \vee ... \rightarrow Q$ is to invite the charge that one has failed correctly to identify the property in virtue of which the antecedent necessitates the consequent. Or, to say the same thing the other way around: Someone who offers such a law undertakes a burden to provide positive reason that there isn't a *higher level* but *nondisjunctive* property of things that are $R1 \vee R2...$ in virtue of which they bring it about that Q .

But we still haven't got to the bottom. No doubt, if there is a higher level property that subsumes all the disjuncts of an open disjunction, then we will want to state our laws in terms of it. But *why is that?* Not, surely, because we are prejudiced against disjunctions 'as such'? Rather, I think, it's because, formulas that express closed laws are *stronger* than the corresponding open ones and, *ceteris paribus*, we want to accept the strongest generalizations that our evidence confirms. Accepting the strongest generalizations that one's evidence confirms is *what induction is about*.

I'm suggesting that the intuition that open disjunctions invite closed laws isn't different in kind from the intuition that open *lists* invite universal generalizations. *Swans are white* constrains more worlds (hence supports stronger counterfactuals) than *x is white if (x is swan A, \vee x is swan b, \vee x is swan c, ...)*. So too, and for the same reason, *pain causes avoidance* constrains more worlds than *x causes avoidance if (x is neural state a, \vee x is neural state b... , \vee x is silicon state f... , \vee x is Martian state g, V...)*. The cost of the former generalization is reifying not just the properties of *being-swana* and *being-swana-b* and *being-swana-c...* etc. but also the more abstract, higher level property of *being-a-swana*. The cost of the latter generalization is reifying not just the properties of *being-in-neural-state-a*, and *being-in-neural-state-b...* and *being-in-silicon-state-f...* and *being-in-Martian-state-g*, etc. but also the more abstract,

higher level property of *being-in-pain*. Pretty clearly, standard inductive practice is prepared to hypostatize in both cases. And the success of standard inductive practice suggests pretty clearly that it is right to do so.

It may be, according to this view of the matter, that there are laws about open disjunctive properties after all; God can do whatever He likes, as previously remarked. But at least it's apparent that we have general methodological grounds for preferring a closed law to a corresponding open one, all else equal. Induction is a kind of market prudence: Evidence is expensive, so we should use what we've got to buy all the worlds that we can.

Here's where the discussion stands now:

Kim is wrong about what's wrong with jade. What's wrong with jade is not that it's MR but that it's a (closed) disjunctive property, and closed disjunctive properties are ipso facto not projectible.

Kim is also wrong about the analogy between jade and pain. According to functionalists, pain, but not jade, is MR; viz pain has an *openly* disjunctive realizer. It is to that extent ok for a functionalist to say both that there *aren't* laws about jade and that that there *are* laws about pain.

But a functionalist who says this still needs to explain why we should (why we do) prefer higher level closed laws (*pain leads to avoidance*) to lower level open laws (*states that are $R1 \vee R2 \vee \dots$ lead to avoidance*), all else equal. Why are we prepared to buy closed laws at the cost of reifying high level properties? My story is that this policy complies with an injunction that all of our inductive practice illustrates: *Prefer the strongest claim compatible with the evidence, all else equal*. Quantification over instances is one aspect of rational compliance with this injunction; reification of high level kinds is another.

So, then, everything is fine and all the mysteries—except, of course, the ones about induction itself—have dissolved? Not! (As I'm horrified to hear that they say in California). I think, in fact, that what's really bugging Kim is indeed a metaphysical mystery about functionalism, and that the discussion we've been having so far hardly touches it. Let me try, in the closing section, to articulate what I take to be the trouble.

Conclusion (molto misterioso)

Kim remarks, at one point, that "...when we think about making projections over pain, very much the same worry should arise about their propriety as did for jade. Consider, a possible law: 'Sharp pains...cause anxiety reactions'. Suppose this generalization has been well confirmed for humans. Should we expect on that basis that it will hold also for Martians whose psychology is implemented

(we assume) by a vastly different mechanism? Not if we accept...that psychological regularities hold, to the extent that they do, in virtue of the causal-nomological regularities at the physical implementation level (15–16)”.

Apparently, Kim wants to get from premises 4 and 5 to the conclusion 6. But, on second thought, an enthymeme would

4. Psychological regularities hold only in virtue of implementation level regularities.
5. Martian pain is implemented by “vastly different” mechanisms than ours.
6. We shouldn’t expect ‘pain’ to be projectible over a mixed population of us and Martians; that is, we shouldn’t expect that Martian-pain will be like our-pain in respects that are thus far unexamined.

seem to have crept into this line of argument. To warrant his inference, Kim also needs some such premise as 7.

7. The behaviors of systems that are “vastly different” at the physical level should be expected not to be reliably similar in respect of their nonphysical (eg. of their higher-level/functional) properties. Such similarities as one finds are accidents (cf. the color of jadeA = the color of jadeB).

But, in respect of premise 7, one has serious reason to hesitate. For one thing, artifacts appear to offer an indefinite number and variety of counterexamples; that’s why references to canopeners, mousetraps, camshafts, calculators and the like bestrew the pages of functionalist philosophy. To make a better mousetrap *is* to devise a new kind of mechanism whose behavior is reliable with respect to the high-level regularity ‘live mouse in, dead mouse out’. A *better* mousetrap is a mechanism that is *even more* reliable with respect to this regularity than mousetraps used to be. If it weren’t possible, at least *sometimes*, for quite heterogeneous mechanisms to be reliable in respect of the functional descriptions that they converge upon, new kinds of mousetraps would never work; ingenuity would fail, and mousetrap makers would be doomed to careers of fruitless self-quotation. It looks, in short, like (7) might do as a rule of thumb, but it can’t be anything like a natural necessity.

Likewise outside the realm of artifacts. The very *existence* of the special sciences testifies to reliable macrolevel regularities that are realized by mechanisms whose physical substance is quite typically heterogeneous. Does anybody really doubt that mountains are made of all sorts of stuff? Does anybody really think that, since they are, generalizations about mountains-as-such won’t continue to serve geology in good stead? Damn near everything we know about the world suggests that unimaginably complicated to-ings and fro-ings of bits and pieces at the extreme *microlevel* manage somehow to converge on stable *macro-level* properties.

On the other hand, the ‘somehow’ really is entirely mysterious, and my guess is that that is what is bugging Kim. He just doesn’t see why there should be (how there could be) macrolevel regularities *at all* in a world where, by common consent, macrolevel stabilities have to supervene on a buzzing, blooming confusion of microlevel interactions. Or rather, he doesn’t see why there should be (how there could be) unless, at a minimum, macrolevel kinds are *homogeneous* in respect of their microlevel constitution.⁵ Which, however, functionalists in psychology, biology, geology and elsewhere, keep claiming that they typically aren’t.

So, then, *why is there anything except physics?* That, I think, is what is *really* bugging Kim.⁶ Well, I admit that I don’t know why. I don’t even know how to *think about* why. I expect to figure out why there is anything except physics the day before I figure out why there is anything at all, another (and, presumably, related) metaphysical conundrum that I find perplexing.

I admit, too, that it’s embarrassing for a professional philosopher—a paid up member of the APA, Eastern Division, I assure you—to know as little as I do about why there is macrostructural regularity instead of just physical regularity. I would therefore quite like to take Kim’s way out and dissolve the mystery by denying the premise. Kim wants to *just stipulate* that the only kinds there are are (what he calls) *local*; viz the only kinds there are are the kinds of kinds whose realizers are physically homogeneous. “...the present view doesn’t take away species-restricted mental properties... , although it takes away pain ‘as such’ (25).” More precisely, it doesn’t take away species specific mental properties unless it turns out that they are MR too; if they are, then of course it does.

In effect, Kim wants to make it true *by fiat* that the only projectible kinds are physically homogeneous ones. That’s tempting, to be sure; for then the question why there are *macrolevel* regularities gets exactly the same answer as the question why there are *microlevel* regularities: *Both follow from physical laws*, where a physical law is, by definition, a law that has physical kinds on both ends.

But, for better or worse, you don’t get to decide this sort of thing by fiat; just as you don’t get to avoid the puzzle about why there’s something instead of nothing by stipulating that there isn’t. Only God gets to decide whether there is anything, and, likewise, only God gets to decide whether there are laws about pains; or whether, if there are, the pains that the laws are about are MR. Kim’s picture seems to be of the philosopher impartially weighing the rival claims of empirical generality and ontological transparency, and serenely deciding in favor of the latter. *But that picture won’t do*. Here, for once, metaphysics actually matters⁷, so philosophers don’t get to choose.

Science postulates the kinds that it needs in order to formulate the most powerful generalizations that its evidence will support. If you want to attack the kinds, you have to attack the generalizations. If you want to attack the generalizations, you have to attack the evidence that confirms them. If you want to attack the evidence that confirms them, you have to show that the predictions that the generalizations entail don’t come out true. If you want to show that the predic-

tions that the generalizations entail don't come out true, you have actually to *do the science*. Merely complaining that the generalizations that the evidence support imply a philosophically inconvenient taxonomy of kinds *cuts no ice at all*. So far, anyhow, when the guys in the laboratories actually do the science, they keep finding that mental kinds are typically MR, but that the predictions that intentional psychology entails are, all the same, quite frequently confirmed. Lots of different sorts of micro-interactions manage, somehow or other, to converge on much the same macro-stabilities. The world, it seems, runs in parallel, at many levels of description. You may find that perplexing; you certainly aren't obliged to like it. But I do think we had all better learn to live with it.

Notes

*Many thanks to Ned Block, Howard Stein, Ernie Lepore, Joe Levine, and especially Barry Loewer, for helpful criticism of an earlier draft.

1. From here on, I'll honor the distinction between psychological (etc.) states and psychological (etc.) terms only where it matters. For much of the discussion we're about to have, it doesn't. Roughly, psychological states are what the terms in psychological theories denote if the theories are true.
2. 'Real' clearly can't mean exceptionless in this context. But what Kim has against psychological laws isn't their failure to be strict. (Here again, contrast Davidson (1980); see also Fodor, (1991) and Schiffer (1991).)
3. More precisely: If jade is nephrite or jadite, then it's metaphysically necessary that jade is nephrite or jadite. I don't want to speculate on what we would (/should) do if, for example, we were to find that there is yet a third kind of stuff in our jade samples.
4. I'm neutral on the hard question whether a property can have a metaphysically possible realizer that isn't one of its nomologically possible realizers.
5. Note that what's mysterious isn't macrostructure per se; it's unreducible (autonomous) macrostructure. When macrokinds are metaphysically identical to microkinds, laws about the latter imply laws about the former; likewise when macroregularities are logical or mathematical constructions out of microregularities, as in the "Game of Life" described in Dennett (1991). Pace Dennett, such cases do not illuminate (what functionalists take to be) the metaphysical situation in the special sciences. To repeat: autonomy implies 'real' (viz projectible) patterns *without reduction*.
6. To be sure, Kim is also bugged about problems of causal and explanatory overdetermination (see Kim, 1993, *passim*.) But these are plausibly just the other side of the metaphysical problem about levels. Microlevel properties are projectible by consensus. If macrolevel properties can be both projectible and autonomous, it looks like a given causal transaction could be irreducibly 'covered' by more than one causal law; and, presumably reference to any of the covering laws could constitute a causal explanation of the transaction. That, I suppose, is what the issues about causal and explanatory overdetermination amount to.
7. And not just for psychology; the parallels in the current evolutionary wars are positively eerie. "...the ultra-Darwinians reveal a thoroughgoing reductionist stance. [They] simply wave at large-scale systems, but only address the dynamics of gene-frequency changes as they see them, arising from competitive reproductive struggle... . Natural-

ists, in contrast, are attuned to the hierarchical structure of biological systems. They are convinced that there are processes relevant to understanding evolution that go on within each of these levels—from genes, right on up through populations, species, and ecosystems.(p.5)” “The implacable stability of species in the face of all that genetic ferment is a marvelous demonstration that large-scale systems exhibit behaviors that do not mirror exactly the events and processes taking place among their parts... . Events and processes acting at any one level cannot possibly explain all phenomena at higher levels.(p. 175).” (Eldredge, 1995. A super book, by the way, which everybody ought to read.)

References

- Block, N and Fodor, J. (1972) “What psychological states are not,” *Philosophical Review*, 81, 159–181.
- Block, N. (1978) “Troubles with functionalism,” in C. W. Savage ed., *Minnesota Studies in the Philosophy of Science, Vol IX*. U. of Minnesota Press, Minneapolis, Minn.
- Davidson, D. (1980), “Mental events” in *Essays on Actions and Events*, Clarendon Press, Oxford.
- Dennett, D. (1991), “Real Patterns,” *Journal of Philosophy*, 87, 27–51.
- Eldredge, N. (1995) *Reinventing Darwin*, John Wiley, NY.
- Fodor, J. (1974) “Special sciences,” *Synthese*, 28, 97–115.
- Fodor, J. (1991) “You can fool all of the people some of the time, everything else being equal; hedged laws in psychological explanations,” *Mind*, 100, 1, 19–34.
- Kim, J. (1992) “Multiple realization and the metaphysics of reduction,” *Philosophy and Phenomenological Research* 52, pp. 1–26.
- Kim, J. (1993), *Supervenience and Mind*, Cambridge Studies in Philosophy, Cambridge University Press, England.
- Putnam, H. (1967) “The nature of mental states,” in Rosenthal, ed. qv.
- Rosenthal, D. (ed.) (1991) *The Nature of Mind*, Oxford University Press, NY.
- Schiffer, S. (1991) “Ceteris paribus laws,” *Mind*, 100, 1, 1–17.