

Who's Afraid of Disjunctive Properties?

Louise Antony

A Priori

www.apriori.canterbury.ac.nz

Volume 2

© 2003 Louise Antony

Jaegwon Kim's well-known "explanatory exclusion" argument is usually discussed in the context of the mental causation debate. But I also see in this argument a fundamental challenge to a popular view of the relation between the mental and the physical, namely, the view that mental properties are higher-order properties realized by physical properties. (I'll refer to this view simply as "*Multiple Realizability*" or "MR.") MR aims to maneuver between the Scylla of dualism and the Charybdis of type reductionism by asserting two things about mental properties: first, that they are realized by physical properties,¹ and second, that they are *multiply realizable* by physical properties. The first thing explains how mental properties can be causally efficacious: because a mental property is instantiated via the instantiation of a physical property, the mental property "inherits" all the causal powers associated with the realizing physical property. The second thing certifies the autonomy of the mental: because one and the same mental property can be realized in a variety of different physical properties, the mental property cannot be identified with, and therefore cannot be reduced to, any particular physical realizer property. For purposes of preserving the autonomy of the mental, the "able" in "multiply realizable" is crucial. Even if it were to turn out that mental properties were *actually* co-extensive with specific physical properties, the fact that they were not *necessarily* co-extensive would be enough to show that the mental properties and these physical realizer properties were not the same.

Many philosophers find MR intuitively plausible. Indeed, the ability to countenance creatures with mental lives who happen not to be constituted like us has long been considered an important desideratum for

¹ Hilary Putnam, one of the original proponents of this view, held that mental properties needn't be realized in physical stuff; they could be realized, he insisted, in "soul stuff." I confess I never really understood what he had in mind. Perhaps he was simply trying to make the point that the *functionalism* in the view of mind he was developing did not commit him to materialism. It would then be left open whether the notion of "causal relations" employed in functional specifications somehow carried the requirement the relata were material states or events. To settle that question, we'd need to engage such vexed questions as what counts as "material," and whether the non-material can be involved in causal relations. Fortunately, we needn't take up such questions here, for in this paper, I am specifically interested, in defending *materialist* versions of the thesis that mental properties are multiply realized properties.

an adequate theory of mentality. It was an advantage that behaviorism enjoyed over type identity theory, and one that functionalism readily co-opted. What pedagogue among us has not pumped our students' intuitions in favor of one or the other of these views by bidding them to consider the android Data of Star Trek? In a properly run classroom, all will agree that you don't need a human brain to possess a mind.² Significantly, even John Searle, who maintains that intentionality is a causal power of *brains*, insists that his view leaves open the possibility of artificial or alien minds – a mind needn't be made out of actual brain stuff, he says, as long as the alternate stuff has the same causal powers as brains. (Searle 1992)

But all along, there have been dissonant voices. David Lewis, in his "Mad Pain and Martian Pain" (Lewis 19??) pointed out that advocates of multiple realizability might find themselves faced with an uncomfortable choice: between, on the one hand, honoring the idea that similarity of stuff (brains, nerves, muscles, bones – your standard issue terrestrial components) suffices for similarity of mental states, and on the other, retaining the idea that functional profile is the deciding factor, damn the stuff. Ned Block, in "Troubles with Functionalism," dares us to uphold our functionalist principles when faced with "minds" made out of the citizens of China. More recently, Ruth Millikan has argued that mentality can only be treated scientifically if it is considered to be a *biological* property, and Noam Chomsky (*et tu*, Noam?) seems to agree, averring that "there is no interest in taking 'mental types' to be non-biological, any more than there would be in defining 'chemical' or 'optical types' that share some properties of chemical and optical aspects of the world." (Chomsky 2003, 261). For both Millikan and Chomsky, the proper subject of a serious science is mentality as we find it, instantiated in brains. In Millikan's view, the question whether Data has a mind is as idle a question as whether Emma Bovary had a freckle on her left knee; in Chomsky's view, it's akin to the

² As do the writers of the television series *Star Trek: The Next Generation*. In an episode called "The Measure of Man," a legal hearing is convened to determine whether Data is anything other than a machine, a piece of property. Captain Picard is tapped to "represent" Data (one wonders why the fact that Data was offered "representation" didn't settle the matter at once) and argues that because he is sentient, and has experienced love, he is, in fact, "new life." (You can read all about it at <http://www.cinerhama.com/tvpage/stng/episodes2.htm>) Score one for functionalism.

question of whether airplanes "really" fly – a matter for stipulation, not discovery.

Kim is no doubt sympathetic to all these points. But his particular challenge to MR has more bite. None of the critics mentioned above claims that there's anything *incoherent* about MR. Kim's objection, however, if correct, would show that the arguments standardly given in favor of multiple realizability are self-undermining. I don't think that this aspect of Kim's argument has been widely appreciated, so let me start by showing how this radical challenge emerges from the explanatory exclusion argument.

Kim's first premise is the *Principle of the Causal Closure of the Physical* (CC) – the principle that every physical event (that has a cause at all)³ has a physical cause. It follows immediately from CC that for every physical event, there is a complete causal *explanation* in terms of purely physical properties. Now, Kim says, consider a particular physical event – say, an upward motion of my left forefinger at time *t*. Let us suppose that this particular physical event has a mental cause – say, my wanting to bring my E♭ into better tune. CC requires that this motion of my forefinger must have, in addition to the causal explanation in terms of my desire to play in better tune, a complete physicalistic explanation. But now, Kim argues, it looks like there is a surfeit of causal responsibility. Given the completeness of the physicalistic causal explanation, there is no room for the desire cited in the mentalistic explanation as a "mental cause" to add anything to the causal picture. The physical conditions specified in the physicalistic causal explanation are, by hypothesis, strictly sufficient for bringing about the upward motion of my finger – there's nothing left for my desire to do. Thus, if the mental factor is really distinct from the physical factor, then it must be causally otiose. Equivalently, if the mental factor is causally efficacious, then it must be identical with the physical factor. Either way the defender of the autonomy of the mental loses: either the mental is not autonomous from the physical, or it is epiphenomenal, causally inert.

Now an independent case can be made, alas, for the claim that my desire to play in tune is generally causally inefficacious. But that's hardly the point. Realists about the mental want it to turn out that mental events *can* cause things to happen. Indeed, the best argument for being a realist about the mental is the fact that mental factors like beliefs, desires, pains, and pleasures seem to be so obviously causally responsible for so much of

³ Subsequently, for ease of exposition, I'll omit this qualification.

what we do and so much of what happens to us. So the horn of Kim's dilemma that offers autonomy at the expense of causal efficacy seems like a bad deal. Kim's preferred solution is to give up the autonomy of the mental and embrace type reductionism: mental properties just *are* physical properties.

But this solution can't be adequate if mental properties are multiply realizable. Multiple realizability has *ontological* consequences. Clearly, a property *P* cannot be identical with a property *Q* if there can be instances of *P* that are not instances of *Q*. But to say that a property *M* is multiply realizable is to imply that there are at least two physical realizer properties *P*₁ and *P*₂ such that a) *P*₁ and *P*₂ cannot be co-instantiated, and b) there is a possible world in which *M* is co-instantiated with (and realized by) *P*₁, and there is a possible world in which *M* is co-instantiated with (and realized by) *P*₂. So there can be instances of *M* that are not instances of *P*₁ and instances of *M* that are not instances of *P*₂, and *M* cannot be identical with either *P*₁ or *P*₂.

Kim's response to this ontological argument is diabolical: if you try to evade the reductionist conclusion of the exclusion argument by appeal to multiple realizability, you'll lose even as you win – for to save the *autonomy* of the mental, you'll have to concede its *nomicity*. Notice, Kim says, that for *M* to be genuinely *multiply* realizable, *P*₁ and *P*₂ must be distinct physical kinds – they must possess distinct sets of causal powers. Multiple realizability, then, entails *causal heterogeneity*: every mental property, on the MR account, is associated with a heterogeneous set of causal powers. But in that case, mental “properties” – whatever else they are – *cannot be natural kind properties*. If there is no uniform causal profile associated with a mental property, then there is no basis for generalizing from one instance of a mental property to any other. There could not be, in other words, any *mentalist laws*. It looks like we save *mentality* only to lose *psychology*.

As Kim puts the point:

[M]any philosophers want to argue that [a mental property] *E* is an irreducible property that nonetheless can be a property playing an important role in a special, “higher-level” science. I believe, however, that this position cannot be sustained. For if the “multiplicity” ...of realizers means anything, it must mean that these realizers are causally and nomologically diverse.... All this points to the inescapable conclusion that *E*, because of its causal/nomic heterogeneity, is unfit to figure in laws, and is

thereby disqualified as a useful scientific property. On this approach, then, one could protect *E* but not as a property with a role in scientific laws and explanations. You could insist on the genuine propertyhood of *E* as much as you like, but the victory would be empty. [Kim, 1997, pp. 17-8]

There's a deep irony in Kim's charging that multiply realized properties are *ipso facto* unfit for scientific taxonomies, because it was supposed to be the asystematicity of the class of physical *realizer* properties that provided the best arguments *for* countenancing higher-order properties. Consider, for example, Fodor's argument, in his influential paper, “Special Sciences,” for admitting such special science kinds as “monetary exchange” into one's ontology:

I am willing to believe that physics is general *in the sense that it implies that any event which consists of a monetary exchange* (hence any event which falls under Gresham's Law) *has a true description in the vocabulary of physics and in virtue of which it falls under the laws of physics*. But banal considerations suggest that a description which covers all such events must be wildly disjunctive....What are the chances that a disjunction of physical predicates which covers all these events...expresses a physical natural kind? In particular, what are the chances that such a predicate forms the antecedent or consequent of some proper law of physics? (Fodor 1974 in Moser & Trout, 55-6) [Emphases original]

Fodor's skepticism that “wildly disjunctive” descriptions will figure in laws at the physical level presumably derives from the causal heterogeneity of the items satisfying the description. He continues:

The point is that monetary exchanges have interesting things in common; Gresham's Law, if true, says what one of these interesting things is. But what is interesting about monetary exchanges is surely not their commonalities under *physical* description. A natural kind like a monetary exchange *could* turn out to be co-extensive with a physical natural kind; but if it did, that would be an accident on a cosmic scale. (Fodor 1974 in Moser & Trout, 56).

But now Kim's point comes to the fore: if the set of physical realizers of the economic property MONETARY EXCHANGE is wildly disjunctive and anomic when described in physicalistic terms, that same set cannot

somehow transform itself into a unified and nomic kind simply by being *described* in a new way.

I think a good way to see Kim's point is to reflect on the familiar diagram introduced by Fodor to represent the metaphysical structure of multiple realizability.⁴



The \mathbf{P}_x 's and \mathbf{P}^*_x 's express causally heterogeneous lower-order physical properties, and the \mathbf{M}_x 's express the higher-order properties adverted to in the laws of the special sciences. Fodor argues that the predicates constructed by disjoining the \mathbf{P}_x 's and the \mathbf{P}^*_x 's would not figure in laws of physics, and hence would not express nomic properties. The \mathbf{M}_x 's, on the other hand, capture (at least some of) the interesting things that all the \mathbf{P}_x 's and the \mathbf{P}^*_x 's actually have in common, and thus are fitted for the expression of higher-order regularities.

What Kim is pointing out is that all this has a whiff of magic about it. Either \mathbf{P}_1 's and \mathbf{P}_2 's and \mathbf{P}_3 's (and so on) have something in common or they don't. How can the mere introduction of the new predicate, \mathbf{M}_1 , alter the situation? Fodor's argument does seem to presume that the "kindhood" or nomicity of a *grouping* of objects or events turns on what *vocabulary* we use to pick the grouping out, so lawfulness is relative. He says, after all, that the problem with treating the class of monetary exchanges as a physical natural kind is that it is unlikely that a disjunctive physical "*predicate* forms the antecedent or consequent of some proper law of physics." [my emphasis] The assumption seems to be that if a grouping can be designated by means of a primitive term in the proprietary vocabulary of some particular science, then the grouping counts as a natural kind, *relative to that science*, otherwise not. If so, then one and the same set of objects or events can be nomic and non-nomic, relative to different sciences. Similarly for laws: in order for a generalization to count as a law, its antecedents and consequents must utilize the proprietary vocabulary of some particular science. But generalizations that are not lawful relative to one science, can turn out to be lawful relative to a different science,

provided the terms of the first generalization are necessarily co-extensive with primitive terms of the proprietary vocabulary of the second one.

This way of thinking of nomicity, as relative to mode of description, is at odds with the idea that science taxonomizes according to causal powers (an idea that Fodor himself endorses elsewhere (Fodor, 1987, Ch. 2)). The situation is not improved by the willingness of MR defenders to concede that higher-order properties have no new causal powers of their own, but simply "inherit," on each occasion of their instantiation, whatever cluster of causal powers is possessed by their lower-order realizer on that occasion. If that's the case, then what objectively unifies the items that fall under the higher-order predicate? Why is such a predicate anything more than a typographical shortcut – an abbreviation – of a complicated lower-order disjunction?

Now perhaps the defender of MR will be inclined to object that higher-order properties *do* have new causal powers associated with them. The property \mathbf{M}_1 , for example, is associated with the power to cause \mathbf{M}_2 's. But this response can't allay the current concern. If the physical realizer properties are, by hypothesis, causally heterogeneous, then it's hard to see what this allegedly common causal power could be. The power of \mathbf{M}_1 's to cause \mathbf{M}_2 's looks a *bogus* causal power – an artifact of a decision to designate things in one way rather than another. To see this, look again at Fodor's diagram.



Now suppose we were to choose, for our \mathbf{P} 's and \mathbf{P}^* 's, an *arbitrary* set of physical predicate pairs ($\mathbf{P}_x, \mathbf{P}^*_x$) such that \mathbf{P}_1 's cause \mathbf{P}^*_1 's, \mathbf{P}_2 's cause \mathbf{P}^*_2 's and so forth. Make these as dissimilar as possible: one pair of predicates could apply to rock-hitting-window events and window-breaking events, another to person-taking-Benadryl events and person-falling-asleep events, and so forth. Now let's define two "higher-order" predicates: \mathbf{e} and \mathbf{e} as follows: an item is in the extension of \mathbf{e} just in case it is in the extension of one of the \mathbf{P}_i 's and an item is in the extension of \mathbf{e} just in case it is in the extension of one of the \mathbf{P}^*_i 's. Now it will be true both that \mathbf{e} is distinct from any particular \mathbf{P}_x and that \mathbf{e} 's all have the "novel" power of causing \mathbf{e} 's. But this is ridiculous: \mathbf{e} and \mathbf{e} are bogus, cooked-up properties, and the

⁴ For uniformity of exposition, I replace Fodor's variables " \mathbf{S}_1 " and \mathbf{S}_2 " with " \mathbf{M}_1 " and " \mathbf{M}_2 ," and eliminate his object variables.

power of causing Φ 's is a bogus, cooked-up causal power. And yet nothing said so far about the structure of multiple realization distinguishes this situation from the one that is supposed to obtain when allegedly real and autonomous higher-order properties are multiply realized in diverse physical properties.

Let's recap: the explanatory exclusion argument is supposed to show that either mental properties are identical to physical properties or they are causally inert. The defender of MR retorts that because mental properties are multiply realizable, they simply *are not* identical with any physical realizer property. Kim responds that if that's true, then MR properties can be no more nomic than the disjunction of their realizers, and so cannot be properties that can figure in scientific laws and explanations. The mere existence of a mentalistic *predicate* that in fact applies to a disorderly collection of things of various types hardly secures the reality of any allegedly corresponding property.

But all is not lost for defenders of MR. Notice that there's a region of logical space left unclaimed here. There's an assumption held in common by Kim and his opponents, the defenders of MR, and that is that disjunctive properties – that is, the sorts of properties expressed by disjunctive predicates – are *anomic*. But what happens if that assumption is given up? From Kim's side, what happens is that he loses an objection to the ontological argument for the autonomy of mental properties. So far so good for friends of MR. But this may still look like a pyrrhic victory: it was supposed to be the anomicity of the disjunction of physical realizers that *provided* the ontological argument for the autonomy of mental properties. If the disjunctive properties are themselves nomic, then it looks like there would be no good reason to resist the *identification* of higher-order properties with first-order disjunctive properties, and thus no reason to admit the higher-order properties into our ontology. Again, Fodor makes the argument explicit:

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*. (Fodor 1997, 156) [Emphases original]

But I think this argument goes wrong. Fodor seems to be worried that if we cannot distinguish functional (and more broadly, higher-order) properties from the disjunction of their realizers, then the reductionist will

have won. But what form of "reductionism" can claim this victory? Suppose, for the moment, that we bite the bullet and simply *identify* presumably higher-order mental properties with their lower-order disjunctive associates. If the latter count as "physical" properties, then we would be acceding to the weak reductionist thesis that mental properties are physical properties. But so what? We can still maintain that mental properties are distinct from any of the disjunct properties – we can deny, for example, that mental properties are identical with *neurological* properties – and *that's* the reductionist thesis that Kim is pushing.

Functionalism, in other words, rules out the identification of mental properties with certain *specific* properties that bear a certain *specific* relation to them; it does not rule out the possibility that there be *some* physical property with which mental properties can be identified. What ought to concern the functionalist, in other words, is not whether mental properties turn out, in some way, to *be* physical properties; rather what's crucial is whether mental properties are something over and above *particular* physical realizer properties.

When we look at the matter this way, we can see a connection between the kind of skepticism about MR expressed by Chomsky and Millikan, and the challenge to MR raised by Kim. In all cases, the root doubt is about whether the properties appealed to in characterizations of human psychology will really generalize, as MR says they will, to creatures with very different physical constitutions. What I'm suggesting now is that this focus on the ability of psychology to subsume radically different beings may be obscuring what is to my mind the more fundamental issue, namely whether or not the properties and regularities that figure in the cognitive sciences are the same as or different from those that figure in biology and chemistry, and hence whether the cognitive sciences can legitimately claim autonomy from the others.

Here it is especially instructive to consider Chomsky's position: while he has never deviated from the position that the study of language is the study of a biological entity, the human brain, he has, at the same time, insisted that to study language is to study the brain at a *higher level of abstraction* than is typical for biological inquiry. He is dismissive of recent attempts to provide non-computationalist accounts of language acquisition by "reverse engineering" the brain, working from the bottom, neurological level up. And he is, if anything, even more hostile to teleological approaches that try to infer characteristics of the language faculty either from "proto-linguistic" capacities in near primate relatives, or from

presumed “problem demands” present in the ancestral environment. For Chomsky, there is simply no doubt that the reality of our cognitive lives both requires and permits characterization in terms that are wholly neutral with respect to the particularities of embodiment, and that it is futile to try to understand the details of that embodiment in the absence of adequate descriptions of the functional reality.

In a sense, then, multiple realizability is something of a red herring. What matters, fundamentally, is not whether there could be minds embodied in things other than brains, but rather whether there is a level of reality beyond the level at which brains are normally studied – whether psychological kinds are “really there,” “over and above” the already recognized kinds in chemistry, biology and the other established sciences. If this is what is at stake, then it would not matter if brains turned out to be the only kinds of things that realize minds in any nomologically possible worlds. It would not matter, for example, if it turned out that the only functional description of human mentality rich enough to capture the essential features of our psychological lives was specific in ways that made it nomologically impossible for the functional spec.’s to be satisfied by any kinds of things other than brains. The functional descriptions, and the generalizations given in terms of the psychological categories defined at the functional level would still, in this case, be autonomous from the descriptions, generalizations and categories that turned up at the level of the realizers.

Of course the question will immediately arise what all this talk of psychological kinds’ being “really there” comes to. Indeed – to answer this question would be to give an account of what the *autonomy* of the mental consists in. But difficult as it is, this is the question that we should focus on, not the question of whether there can be non-neurological minds. Of course autonomy is not unrelated to multiple realizability. For all my concessive conjectures above, I must concede that the “autonomy” I’m after does, after all, *entail* multiple realizability, even if the possibilities in question are bare conceptual ones. For I cannot see what one could mean in claiming that psychology is more *abstract* than biology, if one weren’t prepared to say that it’s at least conceptually possible for psychological properties to be instantiated without any particular biological properties being instantiated. What I am suggesting now is that we can attack the question of what makes a category or regularity “real” in a direct way, rather than trying to do it by first establishing the plausibility of multiple realizability. Indeed, as we’ll see when we revisit the argument I developed

above, it’s not going to be possible to make a convincing case for multiple realizability without having something compelling to say about the conditions under which kinds and regularities should count as “real.”

I’ve been intimating that there’s really no reason for defenders of MR to resist the identification of mental properties with disjunctive physical properties, so let me come completely clean. What I really think is that there is no such thing as a “disjunctive property” – rather, there are only disjunctive *predicates*. Moreover, I think that properties are, in themselves, neither “mental” nor “physical” – there are only mentalistic and physicalist predicates, which may or may not express real properties. The classification of predicates depends on the use we make of them: mentalistic predicates are the predicates that occur in mentalistic discourse, and physicalistic predicates are the ones that occur in physicalistic discourse. Parts of both bodies of discourse, are, of course, scientific, and the scientific parts have proprietary predicates – predicates which, according to the science in which they are embedded, “carve nature at its joints.” On this way of doing things, reductionism becomes primarily a thesis about the relation between different vocabularies – different ways of carving up the world. The strong reductionist thesis – the interesting one, the one that Kim wants to defend – can then be restated as the thesis that the proprietary predicates of psychology are not co-extensive with – and therefore do not express the same properties as – the proprietary predicates of the lower-order sciences. Denial of this strong reductionist thesis still leaves open the possibility that mentalistic predicates may be co-extensive with some physicalistic predicates – indeed, with disjunctive ones.⁵

I have not yet explained why I think that disjunctive predicates can pick out nomic properties; and of course I have not yet tackled the question of what makes a property real or autonomous. I’ll try to do both things at once.

First, a softening-up consideration: different languages (sometimes) lexicalize different things. In German, if you wish to speak of a single member of the family of domesticated bovinæ, without committing oneself to the beast’s gender, you can, I am told, use the single word “kuh.” But to do it in English, you must circumlocute: since “cow” refers only to female bovids, and “bull” to male, you must say something like “member of

⁵ I first defended the idea that mental properties should be identified with the disjunction of their realizers in Antony 1999. Clapp develops a similar position in Clapp 2001.

the family of domesticated bovinæ”, or if your Latin’s rusty, “cow or bull.” Now surely the mere fact that English speakers have only a (complex) disjunctive term to do what German speakers can do with a primitive term has no bearing on the relation between the *properties* the two terms express. Being a cow-or-bull is just the same thing as being a kuh – both the English and the German expression track the species property of being a member of the family of domesticated *et cetera*. Undoubtedly there are reasons why English insists on gender marking its bovinds when it’s happy to be gender-blind about its cats, dogs and guinea pigs, but the reasons belong to sociolinguistics, not to metaphysics.⁶

This example nicely supports Lenny Clapp’s criticism of one of David Armstrong’s argument against the existence of disjunctive universals. (Clapp 2001) Armstrong argues that

disjunctive properties offend against the principle that a genuine property is identical in its different particulars. Suppose *a* has a property *P* but lacks *Q*, while *b* has *Q* but lacks *P*. It seems laughable to conclude from these premises that *a* and *b* are identical in some respects. Yet both have the “property”, *P* or *Q*. (Armstrong 1978, 20)

But I see nothing at all laughable about the suggestion that *cows and bulls* “are identical in some respects.” They are, in fact, identical in respect of species membership. Any property that can be projected over members of the *Bovinae Bos taurus* – over kühe, in other words – will project just as well over cow-or-bull’s. More generally, as Clapp points out, there is no reason to regard disjunctive properties as “unreal” in cases where the disjuncts have real commonalities.

Armstrong is right that one could not “conclude” merely from the information that this thing is a cow and that thing is a bull, and hence that they are both “cow-or-bull”’s, that they bear a substantive similarity to each other. But why should we expect that we could? There is no reason to think that, in general, analytic relations among designations are reliable indicators of objective relations among the things designated. We can here utilize Kim’s point in reverse: just as you can’t turn an anomic property into a nomic one by concocting a lexically primitive name for it, neither can you turn a nomic property into an anomic one by designating it disjunctively.

⁶ Then there’s the matter of young horses. Young dogs, whether male or female, are “puppies,” and young cats of either sex are “kittens.” But young horses are either “fillies” if female, or “colts” if male.

As Clapp points out, the question whether a disjunctive predicate designates a nomic property turns entirely on whether there is some real resemblance among the things that satisfy the disjuncts. Determinables, like BEING RED, can be expressed as disjuncts of determinants; this is no reason to conclude that BEING RED is not a real or not a nomic property. (Clapp 2002, pp. 125-6)

If I’ve convinced you that there is no reason to think that disjunctive predicates can never express nomic properties, it remains to give some account of when they do and when they don’t. Disjunctive predicates come easy – my bogus property argument was meant to show that not every disjunctive predicate can claim to express a natural kind. So what distinguishes the nomic disjunctives from the non-nomic ones? To answer this question, let me return to Fodor’s argument about monetary exchanges. I pointed out that Fodor seems, in this argument, to be thinking of “kindhood” as relative to vocabulary:

What are the chances that a disjunction of physical predicates...expresses a physical natural kind? In particular, what are the chances that such a predicate forms the antecedent of consequent of some proper law of physics?

A similar theme is sounded in Putnam’s early arguments for admitting higher-order properties into one’s ideology. Citing the robustness of generalizations like “rigid square pegs with 1” sides won’t fit into round holes with 1” diameters, he asserts the methodological necessity of using descriptions like “rigid” and “round hole” instead of (presumably) more precise characterizations of pegs-and-hole interactions in terms of molecular arrangements:

The fact is that we are much more interested in generalizing to other structures which are rigid and have various geometrical relations, than we are in generalizing to *the very next peg that has exactly this molecular structure*, for the very good reason that there is not going to be a next peg that has exactly this molecular structure. So in terms of real life disciplines, real life ways of slicing up scientific problems, the higher level explanation is far more general, which is why it is *explanatory*. (Putnam 1967, p. 127) [Emphases original]

Although Putnam does not say here that the kind “rigid square peg” is *only* a kind when designated in those terms, he does seem to say that it is our *interests* in prediction and explanation that warrant our *treating* rigid square pegs as constituting a kind. Fodor can be read, similarly, as saying that, for

purposes of *human science*, monetary exchanges have to be understood as such, rather than as a collection of physically disparate events. It appears then, that Fodor and Putnam are both making something like the following argument: while there are many different ways of characterizing reality, some ways facilitate human knowledge gathering, and some don't. Physics – and its proprietary vocabulary – provides a system of classification useful for some purposes, but there are other purposes for which it is not as useful. Some groupings and some regularities are only “visible” to us at more abstract levels of apprehension, and since many of these groupings and regularities are of enormous importance to us as human beings, we have crucial need for vocabularies that capture them. Hence, higher-order vocabularies, and the sciences that embrace them, are, for all practical human purposes, ineliminable.

Now this argument may seem to be giving us human beings way too much metaphysical credit. Just because we find it useful to think in terms of some categories rather than others, or to theorize in terms of some vocabularies rather than others – this is supposed to make those categories *real*? Well, yes. Utility, after all, is a matter of predictions working out, of plans being successfully executed. Not just any old way of carving things up is going to get our bridges built or our pneumonia cured or our economy stimulated. Nor, for that matter, will just any old way of carving things up enable us to predict the trajectory of that car to our right on the freeway, or to figure out what someone who is speaking to us wants us to do. Psychological predicates, like macrophysical predicates, geometric predicates, biological predicates, and economic predicates are used by us to state generalizations, predications and explanations that afford us real epistemic power.

Goodman had a name for predicates like this: he called them “projectible.”

So here's a short answer to the question of when disjunctive predicates express nomic properties: when they are necessarily co-extensive with *projectible* predicates. Goodman, recall, was deflationary about projectibility: as he characterized it, it was a property of linguistic entities -- predicates and hypotheses. His position was that we could say all there is to be said regarding the “problem of induction” by reflecting on human linguistic practice. In deciding which of two competing hypotheses – “All emeralds are green” and “All emeralds are grue” – is actually confirmed by a data set, we can and ought to determine which of the two predicates – “green” and “grue” – is more “entrenched” in human language – more used,

more relied upon. (Of course “grue” forfeits this competition for failure to show up.) This sounds, and is, highly anthropocentric, and indeed, Goodman thinks that there is nothing more to the “reality” of natural kinds than their corresponding to classifications that human beings find useful. He does allow that

In a sense, not the word itself but the class it selects is what becomes entrenched, and to speak of the entrenchment of a predicate is to speak elliptically of the entrenchment of the extension of that predicate.

However he continues:

On the other hand, the class becomes entrenched only through the projection of predicates selecting it; entrenchment derives from the use of language. (Goodman 1983)

What I am suggesting, on the contrary, is that the predicates can only *get* entrenched *if the world cooperates*. If monetary exchanges and rigid square pegs *don't* really have something important in common despite their myriad physical dissimilarities, then “monetary exchange” and “rigid square pegs” could not have the roles in our language that they in fact enjoy. The generality that such descriptions afford us would be useless – treating rigid things as relevantly similar to each other would do us no good. It's only if the practice of picking things out *as rigid* or *as square* really does give us a basis for generalizing to other things that we *also* pick out as rigid or as square that we will find any point to lexicalizing these features. In short, there's a strong abductive argument from the *projectibility* of higher-order predicates to the *reality* of the kinds they designate.

If we can assume, then, that the existence of a projectible predicate is presumptive evidence for the existence of a real kind corresponding to it, we can state crisply the condition necessary for a disjunctive predicate to express a nomic property: it must be necessarily co-extensive with a projectible predicate. Fodor's “wildly disjunctive” predicates are not projectible; neither are my bogus predicates, \mathcal{E} and \mathcal{F} . But, I contend, Fodor's predicates – like the one constructed by disjoining each and every physicalistic description that applies to an event that is a monetary exchange – are ones that are necessarily co-extensive with *projectible* predicates – in this example, with the predicate “monetary exchange.” In contrast, there is, as far as we know, no projectible English predicate that is necessarily co-extensive with either \mathcal{E} or \mathcal{F} , and so no reason – again, as far as we know – to grant that the disjunctive predicates that \mathcal{E} and \mathcal{F} abbreviate express nomic properties.

The qualification “as far as we know” is important. I say that necessary co-extensiveness with a projectible property is *sufficient* for a disjunctive predicate to express a nomic property; I deny that it is *necessary*. Obviously, we have not always had projectible predicates to express properties that we now have reason to regard as nomic; *pace* Goodman, the kinds did not suddenly *become* nomic when we thought to group their members together, much less when we invented names for the groupings. Moreover, it seems perfectly possible that there are natural kinds that we will never know about (and hence never speak about), either because the sun burns out before we discover them, or because our inherent epistemic limitations make it impossible for us ever to recognize them.

One caveat before I continue. I say that the existence of a projectible predicative provides abductive support for the existence of a real, nomic kind. But it has been objected⁷ to this that there can be and often are explanations for the entrenchment of a predicate other than the reality of the kinds they (putatively) designate. “Witches” was apparently pretty well-entrenched – by the sociolinguistic measures that Goodman wants us to rely on – for at least part of European-American history, and “angels” apparently still is. But I would not want to concede either to the Puritans or to today’s spiritualists that either witches or angels exist. Now there are two ways to approach the issue. One is to deny that “witches” and “angels” are entrenched in the language. The way to do this, presumably, would be to demonstrate things like this: a) that there really have been very few (or even no) successful predictions made in terms of “witches” or “angel,” or b) that there is very little interpersonal validity to classification of phenomena as “witchy” or “angelic.” The mere fact that lots of people use these words, in short, is not enough for the words to count as “entrenched.” On the other hand, one can pursue a different strategy: concede the entrenchment, but defeat the abductive inference by offering a *better* explanation of it. In the case of “witches,” this can be done – one can find real similarities in the socio-political circumstances of a large group of women who were accused and convicted of witchcraft, and so can explain the utility of “witch” (though it wasn’t very useful for the women who lost their lives) on the basis of a different kind of commonality than the one that the Puritan elders had in mind. The entrenchment of the predicate is then

⁷ And has been, by the members of my audience at Bowling Green State University, virtually in chorus. I hope that, even if my reply doesn’t convince them about my theory, it at least reassures them about my sanity.

explained by the utility of the socio-political grouping, not some supernaturalistic one.

Applying this framework to the debate about the reality of mental properties,⁸ we get a novel way of understanding the difference between *eliminativism* and *reductionism*: The eliminativist about the mental denies that mentalistic predicates are projectible. Such a person disagrees with the mental realist about the robustness of our psychological attributions, denying either that there are any substantive predictions that can be made on the basis of such attributions, or else that the predictions that are so based are successful. The reductionist, on the other, agrees that psychological predicates are projectible, but thinks that they are so only because they track *biological* kinds. Kim, by this criterion, is what we might call a “conditional” reductionist: he thinks that *if* psychological predicates are projectible, it is only *because* they are co-extensive with well-behaved biological predicates. If the eliminativist is right, mental properties aren’t real; if the reductionist (in this sense of “reductionism”) is right, then mental properties are real, but not autonomous.

I will not discuss eliminativism here.⁹ Instead, I want to return to Kim’s challenge. Kim allows that mental properties can be nomic, but *only if* they are *not* multiply realizable. But I’ve been arguing that it is perfectly possible for a multiply realizable property to be nomic, and hence *possible* for mentalistic predicates to be both projectible and to express higher-order properties, despite the causal heterogeneity of the first-order realizers, and the unprojectibility of the equivalent lower-order disjunctive properties.¹⁰ I want to bolster this argument by saying something more about the kinds of regularities in which higher-order nomic properties participate. To do that, let me look back at my bogus higher-order predicates, **C** and **S**. There is an instructively important difference between the way the “properties” putatively expressed by such predicates as these relate to their (putative)

⁸ Or, as I prefer to say, about the putative extensions of mentalistic predicates. In what follows, for ease of exposition, I’ll speak with the vulgar.

⁹ For an excellent presentation of the arguments against eliminativism, see Rey 1997.

¹⁰ In Antony 1999, I explain in detail why the lower-order disjunctive predicates co-extensive with the higher-order predicates of psychology are certain to be spectacularly unprojectible.

first-order realizer properties, and the way legitimate and entrenched higher-order predicates relate to theirs.

All higher-order predicates, I conceded, are necessarily co-extensive with “wildly” disjunctive lower-order predicates. Kim claims that that’s enough to disqualify higher-order properties from being nomic. All such predicates, he argues, are essentially like the predicate “jade.” Jade, as we all know, is not a natural kind; the term “jade” is simply applied indifferently to samples of two distinct mineral kinds, jadeite and nephrite. For this reason, attempts to project observed properties of jadeite onto hitherto unexamined sample of nephrite is fraught with epistemic risk. The situation with (allegedly) multiply realizable properties, like BELIEVES THAT *P* is, according to Kim, perfectly analogous.¹¹ Given that different physical realizers of a given higher-order property can differ enormously with respect to their causal powers, properties observed to obtain in one sample of things with the higher-order property cannot be safely projected onto all others. (Kim, 1993) My bogus predicates, \mathcal{C} and \mathcal{S} , make Kim’s point even more strikingly: one would be foolish indeed to project properties of a rock hurtling through the air onto a pill in the process of being swallowed.

Fodor makes an interesting reply to this argument of Kim’s. (Fodor 1997) The disjunctive predicate “jadeite or nephrite” expresses what Fodor calls a *closed* disjunction: that is, once we are given the two “realizers,” JADEITE and NEPHRITE, the list of possible realizers of JADE is closed. We cannot, in principle, discover new forms of jade, or equivalently, new “realizers” of JADE. The two mineral kinds exhaust the possible “ways” something can be jade, precisely because what jade is is *not* determined by some set of criteria that various different kinds of things might just happen to turn out to satisfy. Fodor thus disputes Kim’s claim that the applicability of the predicate “jade” is determined by a set of observable macrophysical properties. (Kim 1993, p. 24) He points out that a synthetically produced substance that resembled jade in all its observable macroscopic properties would not, for that reason, *be* jade. To *be* jade, something must be jadeite or nephrite – end of story.

¹¹ Kim’s standard example of a higher-order mentalistic property is PAIN, but because it’s controversial whether being in pain is a functional property, and hence a multiply realizable property, is highly controversial, so I prefer to stick with the property of believing (something).

In contrast, Fodor continues, genuine multiply realized properties correspond to “open” disjunctions.¹² We cannot specify in advance the kinds of assemblages that will or will not realize an MR property: it depends entirely on whether the functional specification that defines the higher-order property is satisfied or not. There are thus no *a priori* limits on the kinds of kinds that I might discover or exploit in order to find a novel realization of pain. A robot constructed in my basement – no matter what it’s constructed out of – will count as a *bona fide* believer of *p* as long as it satisfies the relevant functional conditions.

Here’s another way to put Fodor’s point. It’s a mere *accident* that both jadeite and nephrite count as jade; there is nothing that the two mineral kinds *really* have in common. Or better – if there does happen to be something specific that the two mineral kinds have in common (like a particular set of observable macrophysical properties), it is not *in virtue* of those similarities that both jadeite and nephrite are counted as jade. Thus, if we want to check to make sure that something *is* jade, it’s simply not to the point to check observable macro-features – we have to see whether the thing is either jadeite or nephrite. It is, then, in virtue of being either jadeite or nephrite that something is jade, not *vice versa*. I say that BEING JADEITE OR NEPHRITE is thus *ontologically prior* to BEING JADE.

Realizers of (for example) BELIEVES THAT *P*, however, are all realizers of that property in virtue of something that they do all have in common, namely, satisfaction of the functional description that specifically characterizes believing that *p*. The list of believer-realizers is thus open in the sense that we are stand prepared to count in some utterly new thing, provided it resembles other believers in the relevant respects. We needn’t consult a previously constructed list to determine if some new physical assemblage is or is not a believer; something is first determined to be a believer (by virtue of its functional organization as evidenced by its causal interactions with the world); secondarily, then its physicalistic description goes on the realizer list. BEING A BELIEVER THAT *P* is thus ontologically prior to being a human being in such-and-such a neurological state, being a robot in such-and-such an electronic state, and so on for any particular potential realizer.

All of this will be reflected in the fact that generalizations stated in terms of projectible higher-order predicates will have an epistemic profile

¹² I remind the reader that Fodor does not want to identify higher-order properties with their disjunctive lower-order counterparts.

markedly different from those stated in terms of non-projectible “higher-order” predicates, like “jade,” or like my bogus predicates, e and f . I pointed out earlier that a central theme in the classic arguments for multiply realizable properties was that they were needed to “capture” regularities that would otherwise be missed. And I suggested that the problem with this argument as it stood was that the condition threatened to be vacuous: it’s all too easy to cook up “higher-order” predicates in terms of which we can state generalizations inexpressible in any previously existing vocabulary. But we are now in a position to distinguish such bogus “regularities” as would be expressed by “ e ’s cause f ’s” from the real ones. Because e and f are not projectible – because they express closed disjunctions – there is no *realization-independent* way of confirming the generalization e ’s cause f ’s. Because there is nothing that makes something a e other than being on the relevant list, we cannot generalize from the properties of *this* e to those of any other e ’s. In particular, we cannot infer from the fact that something is on the e - list, that it will produce something that happens to be on the f - list. The only “warrant” that we could possibly have for believing that e ’s cause f ’s, short of an exhaustive examination of cases, is the thin analytic fact that being a cause of an f is constitutive of being a e .

When we are dealing with bogus regularities, ascription of the second-order predicate must depend on a prior identification of the instance in question as a “realizer,” and confirmation of the corresponding generalization as a whole must proceed by prior confirmation of a series of first-order causal claims. The situation is markedly different for projectible second-order predicates and real second-order regularities. We do not need to confirm that a creature is in any particular neurological state in order to determine that it believes that p (although, with a good enough realization theory in hand, we could confirm that something is in a particular neurological state by confirming that it believes that p .) But we could not be in a position to project higher-order predicates in this way unless there really were causal powers that all instances of those predicates had in common.

Now Putnam’s argument should remind us that mentalistic predicates are not, by any means, the only higher-order predicates that we human beings are able to apply in realization-independent ways. The existence of higher-order regularities involving predicates outside psychology demonstrates that human beings are perfectly capable of

discerning causal regularities among physically disparate systems. Simply consider artifacts. As Fodor points out

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. (Fodor 1997, p. 160)

Specification of causal roles of any complexity will generally impose non-trivial constraints on the kinds of structures that can realize them, providing the basis for epistemic confidence that there will be a non-trivial set of causal powers that all realizers will share. Ned Block refers to this fact as “the Disney Principle”: “In Walt Disney movies, teacups think and talk, but in the real world, anything that can do those things needs more structure than a teacup.” (Block 199?, p. 120)

If so, then we can expect that genuine higher-order properties will participate in *many* regularities, not just in those that fall out of their functional definitions. In contrast, the *only* “regularities” in which the bogus properties will participate are the ones enforced by the definitional stipulation – in my example, e ’s cause f ’s. *Real* higher-level regularities will not only be stable and recurrent; at least some of them will be *a posteriori*. The properties that participate in real regularities will be reliably connected, in empirically discoverable ways, to *other* stable and recurrent patterns. If all this is true, then real higher-order properties will be associated with distinctive causal profiles after all – distinctive relative to any of the causal profiles of the various realizer properties. The causal profile of any particular realizer property will be too specific for us to identify its powers with those of the higher-order property – this square peg made of plastic will have causal powers that won’t be shared by that one made of wood. What they will share, however, is the causal power to pass through a (suitably sized) square hole. And this is true whatever the walls of the square hole are made of. But the point would remain *even if it turned out that all square pegs happened to be made of the same material*. Simply being made of wood is not relevant to a peg’s ability to pass through holes of certain shapes and sizes. Shape and size (unsurprisingly) are. Thus it’s not the multiple realization that shows the autonomy of the mental, it’s the *reality of the regularities* in which higher-order properties participate.

We can now see that Kim has misled us. He observes that when a higher-order property

M is instantiated on a given occasion by being realized by P , then the causal powers of *this instance* [of the higher-order property] are

identical with (perhaps a subset of) the causal powers of *P*. (Kim 1995, p. 208)

True enough. But Kim continues:

If the causal then *M* in effect contributes nothing new causally, and *M*'s claim to be a new, irreducible property is put in jeopardy. (Kim 1995, p. 209)

But what does it mean to say that “the causal powers of *M* are identical with those of its realization bases?” The causal powers of *M* are not identical with the causal powers of any particular realization base; rather they are identical with the *intersection* of the causal powers of all the realizers. And this can be (and will be in the case of nomic higher-order properties) a substantial set.¹³ Whether or not *M* “contributes anything new” is irrelevant to the reality of *M*. If *M* participates in real regularities, then that's all that's needed to certify, not only *M*'s existence, but indeed, its scientific respectability. And if it's expressible by a projectible predicate, then it's clear that the regularities are discernible in a realization-independent way.

It's true enough that for a property to be truly “multiply” realizable, its realizers must differ from each other substantially in respect of *total causal powers*. But it doesn't follow from this that the realizers can't share a significant *set* of causal powers. In fact, we can see from the case of physical properties that this must be so. SILVER is a paradigmatically physical property, Kim would agree. And yet, instances of SILVER differ widely with respect to their causal powers. Silver coins, in sufficient numbers, have the power to free a candy bar from a vending machine; my silver ring lacks this power, but compensates by being able to encircle my finger. But heterogeneity of causal powers among instances hardly shows that SILVER isn't a perfectly respectable property with its own distinctive set of causal powers. What are they? To find out, we could take the intersection of the causal powers of all things silver.

Now consider the question whether BEING SILVER adds anything to the causal powers of some instance of silver. BEING SILVER is determined by the atomic structure of the molecules that compose the thing in question. But atomic *structure* is an abstract property. So there will be a multitude of particular molecular arrangements alike only in the abstract respect of sharing atomic structure. In any given instance, given the particular molecular assembly that constitutes that instance, the specific

causal powers of that instance are fully determined. What does or could BEING SILVER add? The question is simply a vexed one. The causal powers associated with any properties above the level of fundamental physics all supervene on the properties of the fundamental particles that compose their instances. Higher-order properties are not distinguished in this way from any others above the fundamental level.¹⁴ The question, therefore, should not be “what does this alleged property add?” but rather, is there or is there not a distinctive causal repertoire associated with this alleged property? And that is a test that multiply realized properties can pass.

¹⁴ I am trying to make short work of a complicated matter. Kim introduces the distinction between “higher-order” and “higher-level” properties in Kim 1995. A higher-level property is one that applies to things at a level of aggregation higher than that of fundamental particles; higher-order properties, on the other hand, apply to objects at the same level of aggregation as their lower-level realizer properties. Kim thinks that the causal powers associated with higher-level properties are new, relative to the next lower-level properties, because they are properties of things that are, in a certain sense, new. I am arguing that this consideration is not decisive. Kim develops his point further in Kim 2000.

¹³ Clapp also makes this point. (Clapp 2001)